

Consistent Distribution-Free K -Sample and Independence Tests for Univariate Random Variables

Ruth Heller

*Department of Statistics and Operations Research
Tel Aviv University
Tel Aviv 69978, Israel*

RUHELLER@POST.TAU.AC.IL

Yair Heller

HELLER.YAIR@GMAIL.COM

Shachar Kaufman

Barak Brill

Malka Gorfine

*Department of Statistics and Operations Research
Tel Aviv University
Tel Aviv 69978, Israel*

SHACHARK@POST.TAU.AC.IL

BARAKBRI@MAIL.TAU.AC.IL

GORFINEM@POST.TAU.AC.IL

Editor: Kenji Fukumizu

Abstract

A popular approach for testing if two univariate random variables are statistically independent consists of partitioning the sample space into bins, and evaluating a test statistic on the binned data. The partition size matters, and the optimal partition size is data dependent. While for detecting simple relationships coarse partitions may be best, for detecting complex relationships a great gain in power can be achieved by considering finer partitions. We suggest novel consistent distribution-free tests that are based on summation or maximization aggregation of scores over all partitions of a fixed size. We show that our test statistics based on summation can serve as good estimators of the mutual information. Moreover, we suggest regularized tests that aggregate over all partition sizes, and prove those are consistent too. We provide polynomial-time algorithms, which are critical for computing the suggested test statistics efficiently. We show that the power of the regularized tests is excellent compared to existing tests, and almost as powerful as the tests based on the optimal (yet unknown in practice) partition size, in simulations as well as on a real data example.

Keywords: bivariate distribution, nonparametric test, statistical independence, mutual information, two-sample test, HHG R package

1. Introduction

Testing if two univariate random variables X and Y are independent of one another, given a random paired sample $(x_i, y_i)_{i=1}^N$, is a fundamental and extensively studied problem in statistics. Classical methods have focused on testing linear (Pearson's correlation coefficient) or monotone (Spearman's ρ , Kendall's τ) univariate dependence, and have little power to detect non-monotone relationships. Recently, there has been great interest in developing methods to capture complex dependencies between pairs of random variables. This interest follows from the recognition that in many modern applications, dependencies of interest may not be of simple forms, and therefore the classical methods cannot capture them. Moreover, in modern applications, thousands of variables are measured simultaneously, thus making it impossible to view the scatter-plots of all the potential pairs of variables of interest. For example, Steuer et al. (2002) searched for pairs of genes that are co-dependent, among thousands of genes measured, using the estimated mutual information as a dependence measure. Reshef et al. (2011) searched for any type of relationship, not just linear or monotone, in large data sets from global health, gene expression, major-league baseball, and the human gut microbiota. They proposed a novel criterion which generated much interest but has also been criticised for lacking power by Simon and Tibshirani (2011) and Gorfine et al. (2011), and for other theoretical grounds by Kinney and Atwal (2014).

A special important case is when X is categorical. In this case, the problem reduces to that of testing the equality of distributions, usually referred to as the K -sample problem (where K is the number of categories X can have). Jiang et al. (2014) searched for genes that are differentially expressed across two conditions (i.e., the 2-sample problem), using a novel test that has higher power over traditional methods such as Kolmogorov–Smirnov tests (Darling, 1957).

For modern applications, where all types of dependency are of interest, a desirable property for a test of independence is consistency against any alternative. A consistent test will have power increasing to one as the sample size increases, for any type of dependency between X and Y . Recently, several consistent tests of independence between univariate or multivariate random variables were proposed. Székely et al. (2007) suggested the distance covariance test statistic, that is the distance (in weighted L_2 norm) of the joint empirical characteristic function from the product of marginal characteristic functions. Gretton et al. (2008) and Gretton and Györfi (2010) considered a family of kernel based methods, and Sejdinovic et al. (2013) elegantly showed that the test of Székely et al. (2007) is a kernel based test with a particular choice of kernel. Heller et al. (2013) suggested a permutation test for independence between two random vectors X and Y , which uses as test statistics the sum over all pairs of points $(i, j), i \neq j$, of a score that tests for association between the two binary variables $I\{d(x_i, X) \leq d(x_i, x_j)\}$ and $I\{d(y_i, Y) \leq d(y_i, y_j)\}$, where $I(\cdot)$ is the indicator function and $d(\cdot, \cdot)$ is a distance metric, on the remaining $N - 2$ sample points. Gretton and Györfi (2010) also considered dividing the underlying space into partitions that are refined with increasing sample size. For the K -sample problem, Székely and Rizzo (2004) suggested the energy test. This test was also proposed by Baringhaus and Franz (2004) and mentioned in Sejdinovic et al. (2013) to be related to the maximum mean discrepancy (MMD) test proposed in Gretton et al. (2007) and Gretton et al. (2012a). Harchaoui et al. (2008) adopted the kernel approach of Gretton et al. (2007) and incorporated the covariance into the test statistic by using the kernel Fisher discriminant.

The tests in the previous paragraph are not distribution-free, i.e., the null distribution of the test statistics depends on the marginal distributions of X and Y . Therefore, the significance of these

tests is typically computed by a permutation test. Another option is to try to estimate the null distribution, but this can be very difficult and can lead to over-conservative bounds. For example, Székely et al. (2007) suggested an asymptotically valid null bound for the significance which they showed was too conservative (see Sejdinovic et al. (2013) for suggestions on how to directly estimate the null distribution). The computational burden of applying the permutation tests to a large family of hypotheses may be great. For example, the yeast gene expression data set from Hughes et al. (2000) contained $N = 300$ expression levels for each of 6,325 *Saccharomyces cerevisiae* genes. In order to test each pair of genes for co-expression, it is necessary to account for multiplicity of $M = 2 \times 10^7$ hypotheses. For the permutation tests of Heller et al. (2013) and Székely et al. (2007), the number of permutations required for deriving a p -value that is below $0.05/M$ is therefore of the order of 10^{10} . Since these test statistics are relatively costly to compute for each hypothesis, e.g., $O(N^2)$ in Székely et al. (2007), and $O(N^2 \log N)$ in Heller et al. (2013), applying them to the family of $M = 2 \times 10^7$ hypotheses is computationally very challenging, even with sophisticated re-sampling approaches such as that of Yu et al. (2011). Distribution-free tests have the advantage over non-distribution-free tests, that quantiles of the null distribution of the test statistic can be tabulated once per sample size, and repeating the test on new data for the same sample size will not require recomputing the null distribution. Therefore, the computational cost is only that of computing the test statistic for each of the hypotheses.

We note that for univariate random variables Székely and Rizzo (2009a) considered using the ranks of each random variable instead of the actual values in the test of distance covariance (Székely et al., 2007), resulting in a distribution-free test. Similarly, for the test of Heller et al. (2013) replacing data with ranks results in a distribution-free test. An earlier work by Feuerverger (1993) defined another test based on the empirical characteristic functions for univariate random variables. The test statistic of Feuerverger (1993) was based on a different distance metric of the joint empirical characteristic function from the product of marginal characteristic functions than that of Székely et al. (2007). Moreover, in Feuerverger’s test the X ’s and Y ’s are first replaced by their normal scores, where the normal scores of the X ’s depend on the data only through their ranks among the X ’s, and similarly the normal scores of the Y ’s depend on the data only through their ranks among the Y ’s, thus making this test distribution-free. Pczos et al. (2012) suggested a kernel-based test which uses the MMD after transforming the data to ranks.

A popular approach for developing distribution-free tests of independence considers partitioning the sample space, and evaluating a test statistic on the binned data. A detailed review of distribution-free partition-based tests is provided in Section 1.1 for the independence problem, and in Section 1.2 for the K -sample problem. In 1.3 we describe our goals and the outline of the present paper.

1.1 Review of Distribution-Free Tests of Independence Based on Sample Space Partitions

For detecting any type of dependence between X and Y , the null hypothesis states that X and Y are independent, $H_0 : F_{XY} = F_X F_Y$, where the joint distribution of (X, Y) is denoted by F_{XY} , and the marginal distributions of X and Y , respectively, are denoted by F_X and F_Y . The alternative is that X and Y are dependent, $H_1 : F_{XY} \neq F_X F_Y$.

Figure 1 shows example partitions of the sample space based on the ranked observations, $\text{rank}(Y)$ versus $\text{rank}(X)$, where a $m \times m$ partition is based on $m - 1$ observations. We refer to such partitions as data derived partitions (DDP). The dependence in the data can be captured by many partitions, and some partitions are better than others.

Hoeffding (1948b) suggested a test based on summation of a score over all N 2×2 DDP of the sample space, which is consistent against any form of dependence if the bivariate density is continuous. Hoeffding's test statistic is

$$\iint N \left\{ \hat{F}_{XY}(x, y) - \hat{F}_X(x)\hat{F}_Y(y) \right\}^2 d\hat{F}_{XY}(x, y),$$

where \hat{F} denotes the empirical cumulative distribution function. Blum et al. (1961) showed that Hoeffding's test statistic is asymptotically equivalent to $\sum_{i=1}^N (o_{1,1}^i o_{2,2}^i - o_{1,2}^i o_{2,1}^i)^2 / N^4$, where $o_{u,v}^i$, $u, v \in \{1, 2\}$, is the observed count of cell (u, v) in the 2×2 contingency table defined by the i th observation. Thas and Ottoy (2004) noted that by appropriately normalizing each term in the sum, the test statistic becomes the average of all Pearson statistics for independence applied to the contingency tables that are induced by 2×2 sample space partitions centered about observation $i \in \{1, \dots, N\}$. They proved that the weighted version of Hoeffding's test statistic is still consistent.

Partitioning the sample space into finer partitions than the 2×2 quadrants of the classical tests, based on the observations, was also considered in Thas and Ottoy (2004). They suggested that the average of all Pearson statistics on finer partitions of fixed size $m \times m$ may improve the power, but did not provide a proof that the resulting tests are consistent. They examined in simulations only 3×3 and 4×4 partitions. Reshef et al. (2011) suggested the maximal information coefficient, which is a test statistic based on the maximum over dependence scores taken for partitions of various sizes, after normalization by the partition size, where the purpose of the normalization is equitability rather than power. Since computing the statistic exactly is often infeasible, they resort to a heuristic for selecting which partitions to include. Thus, in practice, their algorithm goes over only a small fraction of the partitions they set out to examine. In Section 4 we show that the power of this test is typically low.

1.2 Review of the K -Sample Problem

As in Section 1.1, we focus on consistent partition-based distribution-free tests. For testing equality of distributions, i.e., for a categorical X , one of the earliest and still very popular distribution-free consistent tests is the Kolmogorov–Smirnov test (Darling, 1957), which is based on the maximum score of all N partitions of the sample space based on an observed data point. Aggregation by summation over all N partitions has been considered by Cramer and von Mises (Darling, 1957), Pettitt (Pettitt, 1976) who constructed a test-statistic of the Anderson and Darling family (Anderson and Darling, 1952), and Scholz and Stephens (1987).

Thas and Ottoy (2007) suggested the following extension of the Anderson–Darling type test. For random samples of size N_1 and $N - N_1$, respectively, from two continuous densities, for a fixed m , they consider all possible partitions into m intervals of the sample space of the univariate continuous random variable. They compute Pearson's chi-square score for the observed versus expected (under the null hypothesis that the two samples come from the same distribution) counts, then aggregate by summation to get their test statistics. A permutation test is applied on the resulting test statistic, since under the null all $\binom{N}{N_1}$ assignments of the group labels are equally likely. They show that the suggested statistic for $m = 2$ is the Anderson–Darling test. They examined in simulations only partitions into $m \leq 4$ intervals.

Jiang et al. (2014) suggested a penalized score that aggregates by maximization the penalized log likelihood ratio statistic for testing equality of distributions, with a penalty for fine partitions.

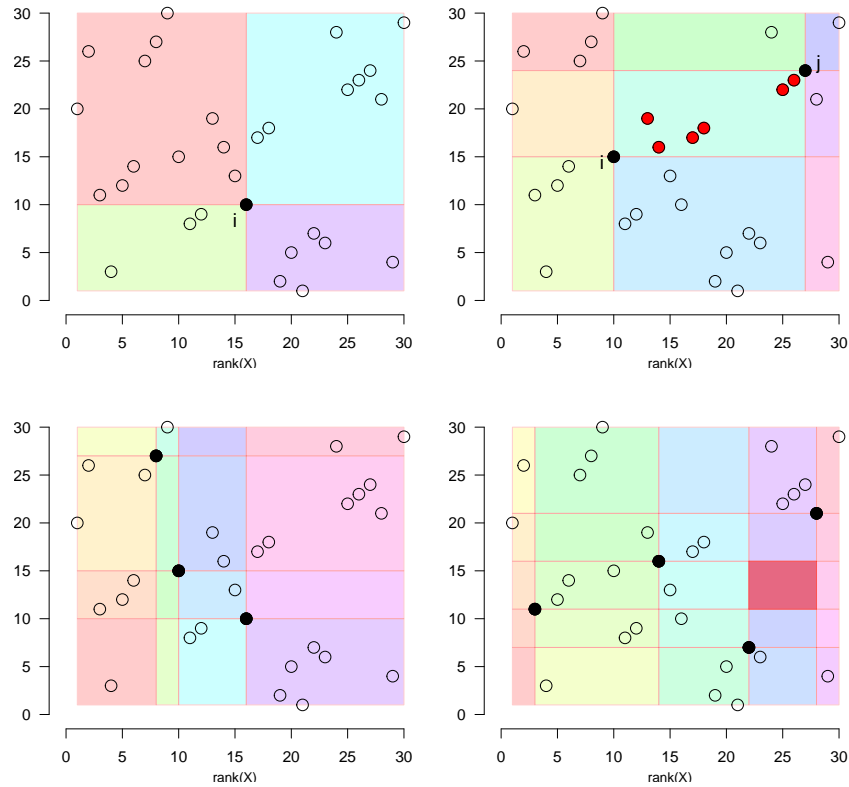


Figure 1: A visualization of the partitioning of the rank–rank plane which is at the basis of the data derived partitions (DDP) tests. Here, $N = 30$, and circles represent observed points. Full black circles represent those observations that were chosen to induce the partition, and different shades represent partition cells. With $m = 2$, all cells are corner cells (top-left); with $m = 3$, the center cell has two vertices which are observed sample points (top-right); with $m = 4$, all internal cells, i.e., cells that are not on the boundary, have at least one observed point vertex (bottom-left); only with $m \geq 5$, there exists at least one internal cell free of observed vertices (bottom-right, internal cell with no vertex which is a sample point is marked in red).

They developed an efficient dynamic programming algorithm to determine the optimal partition, and suggested a distribution-free permutation test to compute the p -value.

Although there are many additional tests for the two-sample problem, the list above contains the most common as well as the most recent developments in this field. Interestingly, when working with ranks, the energy test of Székely and Rizzo (2004) and the Cramer–von Mises test turn out to be equivalent.

1.3 Overview of This Paper

In this work, we suggest several novel distribution-free tests that are based on sample space partitions. The novelty of our approach lies in the fact that we consider aggregation of scores over all partitions of size $m \times m$ (or m for the K -sample case), where m can increase with sample size N , as well as consideration of all ms simultaneously without any assumptions on the underlying distributions. In Section 2 we present the new tests both for the independence problem and for the K -sample problem, with a focus on our regularized scores (that consider all ms) in Section 2.1. We prove that all suggested tests are consistent, including those presented in Thas and Ottoy (2004), and show the connection between our tests and mutual information (MI). In Section 3 we present innovative algorithms for the computation of the tests, which are essential for large m since the computational complexity of the naive algorithm is exponential in m . Simulations are presented in Section 4. Specifically, in Section 4.1 we show that for the two-sample problem for complex distributions there is a clear advantage for fine partitions, while for simple distributions rougher partitions have an advantage. In Section 4.2 we show that, for the independence problem, typically finer partitions have a clear advantage for complex non-monotone relationships. On the other hand, for simpler relationships, there is an advantage for rougher partitions. We further demonstrate the ability of our regularized method (which aggregates over all partitions) to adapt and find the best partition size. Moreover, in simulations we show that for complex relationships all these tests are more powerful than other existing distribution-free tests. In Section 5 we analyze the yeast gene expression data set from Hughes et al. (2000). With our distribution-free tests, we discover interesting non-linear relationships in this data set that could not have been detected by the classical tests, contrary to the conclusion in Steuer et al. (2002) that there are no non-monotone associations. Efficient implementations of all statistics and tests described herein are available in the R package *HHG*, which can be freely downloaded from the Comprehensive R Archive Network, <http://cran.r-project.org/>. Null tables can be downloaded from the first author’s web site.

2. The Proposed Statistics

We assume that Y is a continuous random variable, and that X is either continuous or discrete. We have N independent realizations $(x_1, y_1), \dots, (x_N, y_N)$ from the joint distribution of X and Y . Our test statistics will only depend on the marginal ranks, and therefore are distribution free, i.e., their null distributions are free of the marginal distributions F_X and F_Y .

Test statistics for the K -sample problem We first consider the case that X is categorical with $K \geq 2$ categories. In this case, a test of association is also a K -sample test of equality of distributions. For N observations, there are $\binom{N+1}{2}$ possible cells, and $\binom{N-1}{m-1}$ possible partitions of the observations into m cells, where a cell is an interval on the real line. Since the cell membership of

observations is the same regardless of whether the partition is defined on the original observations or on the ranked observations, and the statistics we suggest only depend on these cell memberships, we describe the proposed test statistics on the ranked observations, $\text{rank}(Y) \in \{1, \dots, N\}$. Let Π_m denote the set of partitions into m cells. For any fixed partition $\mathcal{I} = \{i_1, \dots, i_{m-1}\} \subset \{1.5, \dots, N - 0.5\}$, $i_1 < i_2 < \dots < i_{m-1}$, $\mathcal{C}(\mathcal{I})$ is the set of m cells defined by the partition. For a cell $C \in \mathcal{C}(\mathcal{I})$, let $o_C(g)$ and $e_C(g)$ be the observed and expected counts inside the cell for distribution $g \in \{1, \dots, K\}$, respectively. The expected count $e_C(g)$ is the width of cell C based on ranks multiplied by N_g/N , where N_g is the total number observations from distribution g : $e_{[i_l, i_{l+1}]}(g) = (i_{l+1} - i_l) \times N_g/N$, where $l \in \{0, \dots, m-1\}$, $i_0 = 0.5$ and $i_m = N + 0.5$. We consider either Pearson's score or the likelihood ratio score for a given cell C ,

$$t_C \in \left\{ \sum_{g=1}^K \frac{[o_C(g) - e_C(g)]^2}{e_C(g)}, \sum_{g=1}^K o_C(g) \log \frac{o_C(g)}{e_C(g)} \right\}. \quad (1)$$

For a given partition \mathcal{I} , the score is $T^{\mathcal{I}} = \sum_{C \in \mathcal{C}(\mathcal{I})} t_C$ (where if $t_C = \sum_{g=1}^K o_C(g) \log \frac{o_C(g)}{e_C(g)}$ then $T^{\mathcal{I}}$ is the likelihood ratio given the partition). Our test statistics aggregate over all partitions by summation (Cramer–von Mises-type statistics) and by maximization (Kolmogorov–Smirnov-type statistics):

$$S_m = \sum_{\mathcal{I} \in \Pi_m} T^{\mathcal{I}}, \quad M_m = \max_{\mathcal{I} \in \Pi_m} T^{\mathcal{I}}. \quad (2)$$

Tables of critical values for given sample sizes N_1, \dots, N_K can be obtained for (very) small sample sizes by generating all possible $N!/(\prod_{g=1}^K N_g!)$ reassignments of ranks $\{1, \dots, N\}$ to K groups of sizes N_1, \dots, N_K and computing the test statistic for each reassignment. The p -value is the fraction of reassignments for which the computed test statistics are at least as large as observed. When the number of possible reassignments is large, the null tables are obtained by large scale Monte Carlo simulations (we used $B = 10^6$ replicates for each given sample size N_1, \dots, N_K). For each of the B reassignments selected at random from all possible reassignments, the test statistic is computed. Clearly, the B computations do not depend on the data, hence the tests based on these statistics are distribution free. Again, the p -value is the fraction of reassignments for which the computed test statistics are at least as large as the one observed, but here the fraction is computed out of the $B + 1$ assignments that include the B reassignments selected at random and the one observed assignment, see Chapter 15 in Lehmann and Romano (2005). The test based on each of these statistics is consistent:

Theorem 1 *Let Y be continuous, and X categorical with K categories. Let N_g be the total number of observations from distribution $g \in \{1, \dots, K\}$, and $N = \sum_{g=1}^K N_g$. If the distribution of Y differs at a continuous density point y_0 across values of X in at least two categories, label these 1 and 2, $\lim_{N \rightarrow \infty} \frac{\min(N_1, N_2)}{N} > 0$, and m finite or $\lim_{N \rightarrow \infty} m/N = 0$, then the distribution-free permutation tests based on S_m and M_m are consistent.*

We omit the proof, since it is similar to (yet simpler than) the proof of Theorem 2 below.

Test statistics for the independence problem We now consider the case that X is continuous. For N pairs of observations, there are $\binom{N-1}{m-1} \times \binom{N-1}{m-1}$ partitions of the sample space into $m \times m$

cells, where a cell is a rectangular area in the plane. We refer to these partitions as the all derived partitions (ADP) and denote this set by Π_m^{ADP} . Since the cell membership of observations is the same regardless of whether the partition is defined on the original observations or on the ranked observations, and the statistics we suggest only depend on these cell memberships, we describe the proposed test statistics on the ranked observations, so the N pairs of observations are on the grid $\{1, \dots, N\}^2$. For any fixed partition $\mathcal{I} = \{(i_1, j_1), \dots, (i_{m-1}, j_{m-1})\} \subset \{1.5, \dots, N - 0.5\}^2$, $i_1 < i_2 < \dots < i_{m-1}$, $j_1 < j_2 < \dots < j_{m-1}$, $\mathcal{C}(\mathcal{I})$ is the set of $m \times m$ cells defined by the partition. For a cell $C \in \mathcal{C}(\mathcal{I})$, let o_C and e_C be the observed and expected counts inside the cell, respectively. The expected count in cell C with boundaries $[i_k, i_{k+1}] \times [j_l, j_{l+1}]$ is $e_C = (i_{k+1} - i_k) \times (j_{l+1} - j_l)/N$, where $k, l \in \{0, \dots, m-1\}$, $i_0 = j_0 = 0.5$, $i_m = j_m = N + 0.5$. As with the K -sample problem, we consider either Pearson's score or the likelihood ratio score for a given cell C ,

$$t_C \in \left\{ \frac{(o_C - e_C)^2}{e_C}, o_C \log \frac{o_C}{e_C} \right\}. \quad (3)$$

For a given partition \mathcal{I} , the score is $T^{\mathcal{I}} = \sum_{C \in \mathcal{C}(\mathcal{I})} t_C$ (where if $t_C = o_C \log \frac{o_C}{e_C}$ then $T^{\mathcal{I}}$ is the likelihood ratio given the partition). As above, we consider as test statistics aggregation by summation and by maximization:

$$S_{m \times m}^{ADP} = \sum_{\mathcal{I} \in \Pi_m^{ADP}} T^{\mathcal{I}}, \quad M_{m \times m}^{ADP} = \max_{\mathcal{I} \in \Pi_m^{ADP}} T^{\mathcal{I}}. \quad (4)$$

We consider another test statistic based on DDP, where each set of $m-1$ observed points in their turn define a partition (see Figure 1). This variant has a computational advantage over the ADP statistic for $m < 5$, see Remark 3.1. Since all observations have unique values, the remaining $N - (m-1)$ points are inside cells defined by the partition. There are $\binom{N}{m-1}$ partitions, denote this set of partitions by Π_m^{DDP} . As before, since the cell membership of observations is the same regardless of whether the partition is defined on the original observations or on the ranked observations, and the statistics we suggest only depend on these cell memberships, we describe the proposed test statistics on the ranked observations. For a cell $C \in \mathcal{C}(\mathcal{I})$, where $\mathcal{I} \in \Pi_m^{DDP}$, the boundaries of C are not necessarily defined by two sample points, as depicted at the bottom right panel of Figure 1. We refer to r_l and r_h as the lower and upper values of the ranks of X in C , and to s_l and s_h as the lower and upper values of the ranks of Y in C , where $r_l, r_h, s_l, s_h \in \{1, \dots, N\}$. Let o_C and e_C be the observed and expected counts strictly inside the cell, respectively. The expected count in cell C with rank range $[r_l, r_h] \times [s_l, s_h]$ is $e_C = (r_h - r_l - 1)(s_h - s_l - 1)/[N - (m-1)]$. We consider Pearson's score or the likelihood ratio score for a given cell C , and define t_C as in (3). For a given partition \mathcal{I} , the score is $T^{\mathcal{I}} = \sum_{C \in \mathcal{C}(\mathcal{I})} t_C$, and similarly to (4) we define

$$S_{m \times m}^{DDP} = \sum_{\mathcal{I} \in \Pi_m^{DDP}} T^{\mathcal{I}}, \quad M_{m \times m}^{DDP} = \max_{\mathcal{I} \in \Pi_m^{DDP}} T^{\mathcal{I}}. \quad (5)$$

For each of the test statistics in (4) and (5), tables of exact critical values for a given sample size N can be obtained for small N by generating all possible $N!$ permutations of $\{1, \dots, N\}$. For each permutation $(\pi(1), \dots, \pi(N))$, the test statistic is computed for the reassigned pairs $(1, \pi(1)), \dots, (N, \pi(N))$. Clearly, the computation of these null distributions does not depend on the data, hence the tests based on these statistics are distribution free. As in the case of the K -sample problem, the p -value is the fraction of permutations for which the computed test statistics

are at least as large as the one observed, and when the number of possible permutations is large, the critical values are obtained by large scale Monte Carlo simulations. The test based on each of these statistics is consistent:

Theorem 2 *Let the joint density of X and Y be $h(x, y)$, with marginal densities $f(x)$ and $g(y)$. If there exists a point (x_0, y_0) such that $h(x_0, y_0)$ is continuous and $h(x_0, y_0) \neq f(x_0)g(y_0)$, i.e., there is local dependence at a continuous density point, and if m is finite or $\lim_{N \rightarrow \infty} m/\sqrt{N} = 0$, then the distribution-free permutation tests based on the following test statistics are consistent:*

1. *The test statistics aggregated by summation: $S_{m \times m}^{DDP}$ and $S_{m \times m}^{ADP}$.*
2. *The test statistics aggregated by maximization: $M_{m \times m}^{DDP}$ and $M_{m \times m}^{ADP}$.*

A proof is given in Appendix A. For finite m , as $N \rightarrow \infty$, our test statistics converge to population measures of deviation from independence that are zero if and only if the independence hypothesis is true (see Hoeffding (1948a) for a discussion of a similar population quantity when $m = 2$). For m growing with N , our test statistics estimate the mutual information, as detailed below.

We note that Thas and Ottoy suggested S_m with $t_C = \sum_{g=1}^2 \frac{(o_C(g) - e_C(g))^2}{e_C(g)}$ in Thas and Ottoy (2007), and $S_{m \times m}^{DDP}$ using Pearson's score for finite m in Thas and Ottoy (2004). However, they examined in simulations only $m \leq 4$. Thanks to the efficient algorithms we developed, detailed in Section 3, we are able to test for any $m \leq N$ in the K -sample problem, and for aggregation by summation in the test of independence. If the aggregation is by maximization in the test of independence, the algorithm, detailed in Section 3, is exponential in m and thus the computations are feasible only for $m \leq 4$.

We shall show in Section 4 that the power of the test based on a summation statistic can be different from the power of the test based on a maximization statistic, and which is more powerful depends on the joint distribution. However, for both aggregation methods, using $m > 3$ partitions improves power considerably for complex settings. Therefore, in complex settings our tests with $m > 3$ have a power advantage over the classical distribution-free tests, which focused on rough partitions, typically $m = 2$.

Connection to the MI An attractive feature of the statistics S_m and $S_{m \times m}^{ADP}$, for m large enough, is that they are directly associated with the MI ($I_{XY} = \int h(x, y) \log[h(x, y)/\{f(x)g(y)\}] dx dy$ for continuous X and Y). MI is a useful measure of statistical dependence. The variables X and Y are independent if and only if the MI is zero. Estimated MI is used in many applications to quantify the relationships between variables, see Steuer et al. (2002), Paninski (2003), Kinney and Atwal (2014) and references within. Although many works on MI estimation exist, no single one has been accepted as a state-of-the-art solution in all situations (Kinney and Atwal, 2014). A popular estimator among practitioners due to its simplicity and consistency is the *histogram* estimator, where the data are binned according to some scheme and the empirical mutual information of the resulting partition, i.e, the likelihood ratio score, is computed. Intuitively, one can expect that the statistic $S_{m \times m}^{ADP}$, properly normalized, can also serve as a consistent estimator of the mutual information, when the contingency tables are summarized by the likelihood ratio statistic, since it is the average of histogram estimators, over all partitions. This intuition is true despite the fact that the number of partitions goes to infinity, since we show that the convergence is uniform and that the fraction of "bad" partitions (i.e., partitions with cells that are too big or too small) is small, as long as m goes to infinity at a slow enough rate.

Theorem 3 *Suppose X is categorical with K categories and Y is continuous. Let N_g be the total number of observations from distribution $g \in \{1, \dots, K\}$, and $N = \sum_{g=1}^K N_g$. If $\lim_{N \rightarrow \infty} \frac{N_g}{N} > 0$ for $g = 1, \dots, K$, $\lim_{N \rightarrow \infty} \frac{m}{N} = 0$, and $\lim_{N \rightarrow \infty} m = \infty$, then $\frac{S_m}{N \binom{m-1}{m-1}}$ is a consistent estimator of the MI.*

Theorem 4 *Suppose the bivariate density of (X, Y) is continuous with bounded mutual information. If $\lim_{N \rightarrow \infty} m/\sqrt{N} = 0$, and $\lim_{N \rightarrow \infty} m = \infty$, then $\frac{S_{m \times m}^{ADP}}{N \binom{N-1}{m-1} \times \binom{N-1}{m-1}}$ is a consistent estimator of the MI.*

See Appendix B for a proof of Theorem 4. The proof of Theorem 3 is omitted since it is similar to that of Theorem 4. See Appendix D for a simulated example of MI estimation using $S_{m \times m}^{DDP}$, $S_{m \times m}^{ADP}$, and the histogram estimator. The ADP estimator is the least variable, as is intuitively expected since it is the average over many partitions.

Remark 2.1 *In this work we assume there are no ties among the continuous variables. In our software, tied data are broken randomly, so that our test remains distribution free. An alternative approach, which is no longer distribution free, is a permutation test on the ranks, with average ranks for ties. Then a tied observation, that falls on the border of a contingency table cell, receives equal weight in each of the cells it borders with.*

2.1 The Proposed Regularized Statistics

An important parameter of the statistics proposed above is m , the partition size. A poor choice of m may lead to substantial power loss: if m is too small or too large, it may lack power to discover complex non-monotone relationships. For example, consider the three simulation settings for the two-sample problem in the first row of Figure 2. The best partition for setting 1, “normal vs. normal with delta”, for small sample sizes, is intuitively to divide the real line into three cells: until the start of the narrow peak, the support of the narrow peak, and after the peak ends. Moreover, the best aggregation method is by maximization, not summation, since there are very few good partitions that capture the peak and aggregation by summation using $m = 3$ will aggregate many bad partitions that miss the peak. Therefore, we expect that M_3 will be the most powerful test statistic for setting 1. For setting 2, “Mix. 3 vs. 4 components”, intuitively it seems best to partition into more than seven cells, and that many partitions will work well. For setting 3, “normal vs. normal with many deltas”, it seems best to partition into many cells. Indeed, the power curves in Figure 3 show that for the first setting, M_m is optimal at value $m = 3$, yet if we use this value for the second setting, the test has 20% lower power than optimal power (which is 86% at $m = 10$), and if we use this value for the third setting, the test has 58% less power than the optimal power (which is 88% at $m = 34$). Since the optimal choice of m is unknown in practice, we suggest two types of regularizations which take into consideration the scores from all partition sizes.

The combined p -values statistic The first type of regularization we suggest is to combine the p -values from each m , so that the test statistic becomes the combined p -value. Specifically, let p_m be the p -value from a test statistic based on partition size m , be it S_m or M_m for the K -sample problem, or $S_{m \times m}^{ADP}$ or $S_{m \times m}^{DDP}$ for the independence problem. Due to the computational complexity, we do not consider a regularized score for $M_{m \times m}$. We consider as test statistics the minimum p -value,

$\min_{m \in \{2, \dots, m_{\max}\}} p_m$, as well as the Fisher combined p -value, $-\sum_{m=2}^{m_{\max}} \log p_m$. These combined p -values are not p -values in themselves, but their null distribution can be easily obtained from the null distributions of the test statistics for fixed m s, as follows: (1) for each of B permutations, compute the test statistics for each $m \in \{2, \dots, m_{\max}\}$; (2) compute the p -value of each of the resulting statistics, so for each permutation, we have a set of p -values $p_2, \dots, p_{m_{\max}}$ to combine; (3) combine the p -values for each of the B permutations. Choose B to be large enough for the desired accuracy of approximation of the quantiles of the null distribution of the combined p -values used for testing. Obviously, since the combined p -values are based on the ranks of the data, they are distribution-free. Since they do not require fixing m in advance, they are a practical alternative to the tests that require m as input.

In order to examine how close this regularized score is to the optimal m (i.e., the m with highest power), we looked at the distribution of the m s with minimum p -values in 20,000 data samples from the above-mentioned three simulation settings. For these settings, using the aggregation by maximization statistic, the median m of the minimal p -value was: 3 for the first setting, 9 for the second setting, and 33 for the third setting. Moreover, the first and third quartiles were 3 to 5 for the first setting, 7 to 14 for the second setting, and 19 to 60 for the third setting. We conclude that for these examples, the m that achieves the minimum p -values in most runs was remarkably close to the optimal m (which was 3, 10, and 34 for settings 1, 2, and 3, respectively), suggesting that the power of the minimum p -value statistic is close to that of the statistic with optimal m . Indeed, the power of the minimum p -values in settings 1-3 using aggregation by maximum was 0.825, 0.799, and 0.785, whereas the power using the (unknown in practice) optimal m in settings 1-3 was 0.894, 0.86, and 0.88, respectively. Further empirical investigations detailed in Section 4 give additional support to this regularization method.

An extensive numerical investigation, partially summarized in Appendix G, led us to choose the minimum p -value as the preferred regularization method. Between the two combining functions, we preferred the minimum over Fisher, since Fisher was far more sensitive to the choice of the range of m for combining (see Table 6). This regularized statistic is consistent, as the next theorems show.

Theorem 5 *Let Y be continuous, and X categorical. Let N_g be the total number of observations from distribution $g \in \{1, \dots, K\}$, and $N = \sum_{g=1}^K N_g$. If the distribution of Y differs at a continuous density point y_0 across values of X in at least two categories, label these 1 and 2, $\lim_{N \rightarrow \infty} \frac{\min(N_1, N_2)}{N} > 0$, then the permutation test based on $\min_{m \in \{2, \dots, m_{\max}\}} p_m$ is consistent, if:*

1. *it is based on S_m , $m \in \{2, \dots, m_{\max}\}$, and $\lim_{N \rightarrow \infty} m_{\max}/\sqrt{N} = 0$ or m_{\max} is finite.*
2. *it is based on M_m , $m \in \{2, \dots, m_{\max}\}$ and $\lim_{N \rightarrow \infty} m_{\max}/N = 0$ or m_{\max} is finite.*

Theorem 6 *Let the joint density of X and Y be $h(x, y)$, with marginal densities $f(x)$ and $g(y)$. If there exists a point (x_0, y_0) such that $h(x_0, y_0)$ is continuous and $h(x_0, y_0) \neq f(x_0)g(y_0)$, i.e., there is local dependence at a continuous density point, then the permutation test based on $\min_{m \in \{2, \dots, m_{\max}\}} p_m$ is consistent, if*

1. *it is based on $S_{m \times m}^{DDP}$ or $S_{m \times m}^{ADP}$, and $\lim_{N \rightarrow \infty} m_{\max}/N^{1/3} = 0$ or m_{\max} is finite.*
2. *it is based on $M_{m \times m}^{DDP}$ and $M_{m \times m}^{ADP}$, and $\lim_{N \rightarrow \infty} m_{\max}/\sqrt{N} = 0$ or m_{\max} is finite.*

The proof of Theorem 6 follows in a straightforward way from the proofs of Theorem 2, see Appendix C for details. The proof of Theorem 5 follows similarly from the proof of Theorem 1, and it is omitted.

The penalized statistic The test statistic is the maximum (over all m s) of the statistic plus penalty. For the K -sample problem, Jiang et al. (2014) suggested assigning a prior on the partition scheme and they regularized the likelihood ratio score using this prior. Specifically, they assumed the partition size is Poisson and the conditional distribution on the m partition widths (normalized to sum to one) is *Dirichlet*(1, . . . , 1). This led to their penalty term $-\lambda_0(\log N)(m - 1)$, where $\lambda_0 > 0$ has to be fixed. We assume that the marginal distribution on the partition size is $\pi(m)$ (e.g., Poisson or Binomial), and that the prior probability of selecting \mathcal{I} given m , $\pi(\mathcal{I}|m)$, is uniform. There is an important difference between our uniform discrete prior distribution on partitions of size m and the continuous Dirichlet uniform prior of Jiang et al. (2014). Our prior is uniform on all partitions that truly divide the sample space into m cells, i.e., we cannot have two partition lines between two consecutive samples, since this is actually an $m - 1$ partition. Using the continuous Dirichlet prior results in practice in at most m partitions, but the partition size may also be strictly smaller than m if two partition points lie between two sample points. Therefore, their conditional distribution given the partition size parameter is not necessarily the true size of the partition. Their penalty translates to a conditional probability given a true partition size m of $\frac{(m-1)!}{(N-1)^{(m-1)}}$, compared to our $\pi(\mathcal{I}|m) = 1/\binom{N-1}{m-1}$. Therefore, their score penalizes more severely large m s, and their regularized test statistic has less power when the optimal m is large in our simulations.

For aggregation by maximum in the K -sample problem, we consider the regularized statistic,

$$\max_{m \in \{2, \dots, m_{\max}\}} \{M_m + \log[\pi(\mathcal{I}|m)\pi(m)]\}, \quad (6)$$

where we use the likelihood ratio score per partition. Due to the computational complexity, we do not consider a regularized score for $M_{m \times m}$. For aggregation by summation, our efficient algorithms described in Section 3 enable us to consider the penalized average score per m ,

$$\max_{m \in \{2, \dots, m_{\max}\}} \{SLR_m \pi(\mathcal{I}|m) + \log \pi(m)\} \quad (7)$$

where $SLR_m \pi(\mathcal{I}|m)$ is S_m divided by the number of partitions of size m for the K -sample test, and $S_{m \times m}^{ADP}$ (or $S_{m \times m}^{DDP}$) divided by the number of partitions of size $m \times m$ for the test of independence, using the likelihood ratio score per partition. The null distribution of these regularized statistics is computed by a permutation test, and they are distribution free.

Regularization using priors was less effective than using combined p -values, except when the Poisson prior was used with parameter $\lambda = \sqrt{N}$ (see Table 7). We preferred the first type of regularization since it was at least as effective as regularizing by a Poisson prior, without requiring setting any additional parameters.

3. Efficient Algorithms

For computing the above test statistics for a given N and partition size m , the computational complexity of a naive implementation is exponential in m . We show in Section 3.1 more sophisticated algorithms for computing the aggregation by sum statistics for all $m \in \{2, \dots, N\}$ at once that have complexity $O(N^2)$ for the K -sample problem, and $O(N^4)$ for the independence problem. This is

possible since instead of iterating over partitions, the algorithms iterate over cells. Moreover, the algorithms also enable calculating the regularized sum statistics of Section 2.1 in $O(N^2)$ and $O(N^4)$ for the K -sample and independence problems, respectively, since we just need to go over the list of m scores and for each score S_m check its p-value in our pre-calculated null tables, which requires just an additional $O(N \log(B))$, where B is the null-table size.

We show in Section 3.2 an algorithm with complexity $O(N^3)$ for the K -sample problem for computing the aggregation by maximum for all m at once. This algorithm also enables calculating the regularized maximum statistic of Section 2.1 in $O(N^3)$. The algorithms for aggregating by maximum in the independence problem are exponential in m , and therefore infeasible for modest N and $m > 4$. However, for $m = 3$ and $m = 4$ we provide efficient algorithms with $O(N^2)$ and $O(N^3)$ complexity, respectively, for the DDP test statistics.

3.1 Aggregation by Summation

The algorithms for aggregation by summation are efficient due to two key observations. First, the score per partition is a sum of contributions of individual cells, and the total number of cells is much smaller than the number of partitions (unless $m = 2$ in the K -sample problem, and $m \leq 4$ when using DDP in the independence problem, see Remark 3.1 below). Therefore, we can interchange the order of summation between cells and partitions and thus achieve a big gain in computational efficiency, since it is easy to calculate in how many partitions each cell appears, see equations (8) and (11).

Second, because for a fixed m the number of partitions in which a specific cell appears depends only on the width (and for independence testing, also length) of the cell, the data-dependent computations do not depend on m : the test statistics are the sum of cell scores for every width for the K -sample test, and for every combination of width and length for the independence test, see equations (9) and (12). The complexity of the algorithm remains the same even when the scores are computed for all m s, since the complexity is determined by a preprocessing phase which is shared by all m s. Therefore, the complexity for the regularized scores is the same as the complexity for a single m .

3.1.1 ALGORITHM FOR THE K -SAMPLE PROBLEM

For $g = 1, \dots, K$ (the categories of X) and $r = 1, \dots, N$ (the ranks of Y), we first compute in $O(N)$ A as follows:

$$A(g, r) = \sum_{i=1}^r I(g_i = g),$$

and let $A(g, 0) = 0$. For a cell with rank range $[r_l, r_h]$, where $r_l, r_h \in \{1, \dots, N\}$, using A , the count of observations in category g that fall inside the cell can be computed in $O(1)$ operations as $o_C(g) = A(g, r_h) - A(g, r_l - 1)$. Therefore, for each cell C the contribution of the cell, t_C , can be computed in $O(1)$ time.

Because the score per partition is a sum of contributions of individual cells, S_m is the sum over the score per cell, multiplied by the number of times the cell appears in a partition of size m .

Considering further summing cells of width 1 to N , we may write S_m as follows:

$$S_m = \sum_{\mathcal{I} \in \Pi_m} T^{\mathcal{I}} = \sum_{C \in \mathcal{C}} t_C \sum_{\mathcal{I} \in \Pi_m} I[C \in \mathcal{C}(\mathcal{I})] = \sum_{w=1}^N \sum_{C \in \mathcal{C}(w)} t_C n(w, m, C), \quad (8)$$

where $\mathcal{C}(w)$ is the collection of cells of width w and $n(w, m, C)$ is the number of partitions that include C . For computing $n(w, m, C)$, we differentiate between two possible types of cells: edge cells and internal cells. Edge cells differ from internal cells by having either $r_l = 1$ or $r_h = N$. The number of partitions of order m that include an edge cell of width $w = r_h - r_l + 1$ is given by $\binom{N-1-w}{m-2}$. The number of partitions including a similarly wide internal cell is $\binom{N-2-w}{m-3}$. Therefore, we may write S_m as follows:

$$S_m = \sum_{w=1}^N \binom{N-2-w}{m-3} T_i(w) + \sum_{w=1}^N \binom{N-1-w}{m-2} T_e(w), \quad (9)$$

where $T_i(w) = \sum_{\substack{C \in \mathcal{C}(w), \\ r_l \neq 1 \cap r_h \neq N}} t_C$ and $T_e(w) = \sum_{\substack{C \in \mathcal{C}(w), \\ r_l = 1 \cup r_h = N}} t_C$. The algorithm proceeds as follows.

First, in a preprocessing phase, we calculate $T_i(w)$ and $T_e(w)$ for all $w \in \{1, \dots, N\}$. Since t_C can be calculated in $O(1)$, as described above, the calculation of $T_i(w)$ and $T_e(w)$ for a fixed w takes $O(N)$. Since there are N values for w , we can compute and store all values of $T_i(w)$ and $T_e(w)$ in $O(N^2)$. Also in the preprocessing phase, for all $u, v \in \{0, \dots, N\}$ we calculate and store all $\binom{u}{v}$. This can be done in $O(N^2)$ using Pascal's triangle method.

Given $T_i(w), T_e(w), w = 1, \dots, N-1$ (which are independent of m !), and all $\binom{u}{v}$, we can clearly calculate S_m according to equation (9) for any m in $O(N)$ and therefore for all ms in $O(N^2)$, since $m < N$. Therefore the overall complexity of computing the scores for all ms is $O(N^2)$.

3.1.2 ALGORITHM FOR THE INDEPENDENCE PROBLEM

Let r_i be the rank of x_i among the observed x values, and s_i be the rank of y_i among the y values. The algorithm first computes the empirical cumulative distribution in $O(N^2)$ time and space,

$$A(r, s) = \sum_{i=1}^N I(r_i \leq r \text{ and } s_i \leq s), \quad (r, s) \in \{0, 1, \dots, N\}^2 \quad (10)$$

where $A(0, s) = 0, A(r, 0) = 0$ and $\hat{F}(r, s) = A(r, s)/N$. First, let B be the $(N+1) \times (N+1)$ zero matrix, and initialize to one $B(r_i, s_i)$ for each observation $i = 1, \dots, N$. Next, go over the grid in s -major order, i.e., for every s go over all values of r , and compute:

1. $A(r, s) = B(r, s-1) + B(r-1, s) - B(r-1, s-1) + B(r, s)$, and
2. $B(r, s) = A(r, s)$.

We describe the algorithm for the ADP statistic, which selects partitions on the grid $\{1.5, \dots, N-0.5\}^2$ on ranked data. The main modifications for the DDP statistic are provided in Appendix E. The count of samples inside a cell with rank ranges $r \in [r_l, r_h]$ and $s \in [s_l, s_h]$ can be computed in $O(1)$ operations via the inclusion-exclusion principle:

$$o_C = A(r_h, s_h) - A(r_l - 1, s_h) - A(r_h, s_l - 1) + A(r_l - 1, s_l - 1).$$

Therefore, for each cell C the contribution of the cell t_C can be computed in $O(1)$ time. Because the score per partition is a sum of contributions of individual cells, we may write $S_{m \times m}^{ADP}$ as follows:

$$\sum_{\mathcal{I} \in \Pi_m^{ADP}} T^{\mathcal{I}} = \sum_{C \in \mathcal{C}} t_C \sum_{\mathcal{I} \in \Pi_m^{ADP}} I[C \in \mathcal{C}(\mathcal{I})] = \sum_{w=1}^{N-2} \sum_{l=1}^{N-2} \sum_{C \in \mathcal{C}(w,l)} t_C n(w, l, m, C), \quad (11)$$

where $\mathcal{C}(w, l)$ is the collection of cells of width w and length l and $n(w, l, m, C)$ is the number of partitions that include C . As in the algorithm for the K -sample problem, $n(w, l, m, C)$ depends only on w, l, m , and whether the cell is an internal cell or an edge cell. For simplification, we discuss only the computation of the contribution of internal cells to the sum statistic, and non-internal cells can be handled similarly (as discussed in the algorithm for the K -sample problem). Therefore, our aim is to compute:

$$\sum_{w=1}^{N-2} \sum_{l=1}^{N-2} n(w, l, m) T(w, l), \quad (12)$$

where $T(w, l) = \sum_{C \in \mathcal{C}(w,l)} t_C$ and $n(w, l, m)$ is the number of partitions that include an internal cell of width w and length l when the partition size is m and $\mathcal{C}(w, l)$ is relabelled to be the collection of internal cells of width w and length l .

The algorithm proceeds as follows. First in a preprocessing phase we perform two computations: 1) calculate and store $T(w, l)$ for all pairs $(w, l) \in \{1 \dots N - 2\}^2$. Since t_C can be calculated in $O(1)$, as described above, the calculation of $T(w, l)$ for a fixed (w, l) takes $O(N^2)$ and since there are $(N - 2)^2$ pairs (w, l) the total preprocessing phase takes $O(N^4)$; 2) for all $u, v \in \{0, \dots, N\}$ we calculate and store all $\binom{u}{v}$ in $O(N^2)$ steps using Pascal's triangle method.

Given $T(w, l)$, and all $\binom{u}{v}$, since $n(w, l, m) = \binom{N-2-w}{m-3} \binom{N-2-l}{m-3}$, we can clearly calculate equation (12) for a fixed m in $O(N^2)$ and therefore for all m s in $O(N^3)$. Due to the preprocessing phase the total complexity is $O(N^4)$.

Remark 3.1 *When it is desired to only compute the statistic for very small m , faster alternatives exist. For the two-sample problem, for $m = 2$, the number of partitions is $O(N)$ and therefore an $O(N \log N)$ algorithm can be applied that aggregates over all partitions, and the complexity is dominated by the sorting of the N observations (for $m = 3$, the number of partitions is already $O(N^2)$). Similarly, for the test of independence, the ADP statistic can be calculated in $O(N^2)$ steps for $m = 2$, and the DDP statistic in $O(N^2)$ for $m = 3$, and in $O(N^3)$ for $m = 4$, since this is the order of the number of partitions. Per partition, the computation of the score for $m \leq 4$ is computed in $O(1)$ time since the contribution of a cell can be computed in $O(1)$ time (as shown above). The DDP statistic for $m = 2$ can be computed in $O(N \log N)$, using a similar sorting scheme as that detailed in Heller et al. (2013).*

3.2 Aggregation by Maximization

Algorithm for the K -sample problem Jiang et al. (2014) suggested an elegant and simple dynamic programming algorithm for calculating $\max_m \{M_m - m\lambda(N)\}$ for any function $\lambda(\cdot)$ in $O(N^2)$. We present a modification of their algorithm that enables us to calculate M_m for all m s in $O(N^3)$. As a first step, for all $i \leq N$ and for all $j < i$ we calculate iteratively $M(i, j)$, the maximum score which partitions the first i samples into j partitions. We compute $M(i + 1, j)$ from

$M(a, j - 1)$, $a \leq i$ using:

$$M(i + 1, j) = \max_{a \in \{2, \dots, i\}} \{M(a, j - 1) + t_{[a+0.5, i+1+0.5]}\},$$

where $t_{[a+0.5, i+1+0.5]}$ is the score of the cell from $a+0.5$ to $i+1+0.5$. This calculation takes $O(N)$, and since we have $O(N^2)$ such items to calculate, this step takes $O(N^3)$. Since $M_m = M(N, m)$, the overall complexity for computing the scores for all m s is $O(N^3)$. Note that this algorithm enables us to calculate $\max_{m \in \{2, \dots, m_{\max}\}} \{M_m + \log[\pi(\mathcal{I}|m)\pi(m)]\}$ in $O(N^3)$ for any function $\pi(m)$, thus the regularized test statistic in Section 2.1 can also be computed in $O(N^3)$.

Algorithm for the independence problem The algorithm is the same as described in Remark 3.1 for the ADP statistic for $m = 2$ and for the DDP statistic for $m = 3$ and $m = 4$, with the difference that the aggregation is by maximization (not summation) over the scores per partition. We are not aware of a polynomial-time algorithm for arbitrary m . We discuss ways to reduce the computational complexity in Section 6.

Remark 3.2 *We show in Appendix F that for univariate data the test of Heller et al. (2013) with an arbitrary distance metric, with or without ties, can be computed in $O(N^2)$ in a similar fashion, thus improving their algorithm by a factor of $\log N$ when X and Y are univariate.*

4. Simulations

In simulations, we compared the power of our different test statistics in a wide range of scenarios. All tests were performed at the 0.05 significance level. Look-up tables of the quantiles of the null distributions of the test statistics for a given N were stored. Power was estimated by the fraction of test statistics that were at least as large as the 95th percentile of the null distribution. The null tables were based on 10^6 permutations.

The noise level was chosen separately for each configuration and sample size, so that the power is reasonable for at least some of the variants. This enables a clear comparison using a range of scenarios of interest. Since the power was very similar for the Pearson and likelihood ratio test statistics, only the results of the likelihood ratio test statistic are presented.

The simulations for the two-sample problem are detailed in 4.1, and for the independence problem in 4.2. The analysis was done with the R package *HHG*, now available on CRAN.

4.1 The Two-Sample Problem

We examined the power properties of the statistic aggregated by summation as well as by maximization for $m \in \{2, \dots, N/2\}$, as well as the minimum p -value statistic, $\min_{m \in \{2, \dots, m_{\max}\}} p_m$. We display here the results for $m_{\max} = 149$ and $N = 500$. However, the choice of m_{\max} has little effect on power, see Appendix G for results with other values of m_{\max} . Also, see Appendix G for the results for $m_{\max} = 29$ and $N = 100$.

We compared these tests to six two-sample distribution-free tests suggested in the literature. We compared to Wilcoxon's rank sum test, since it is one of the most widely used tests to detect location shifts. We compared to two consistent tests suggested recently in the literature, the test of Jiang et al. (2014), referred to as DS, and the test of Heller et al. (2013) on ranks, referred to as HHG on ranks. Finally, we compared to the classical consistent tests of Kolmogorov and Smirnov,

referred to as KS, of Cramer and von Mises (which is equivalent to the energy test of Székely and Rizzo (2004) on ranks), referred to as CVM, and of Anderson and Darling, referred to as AD.

We examined the distributions depicted in Figure 2. The three scenarios in the third row were examined in Jiang et al. (2014). The remaining scenarios were chosen to have different numbers of intersections in the densities, ranging from 2 to 18, in order to examine the effect of partition size m on power when the optimal partition size increases, as well as verify that the regularized statistic has good power. The scenarios also differ by the range of support of where the differences in the distributions lie (specifically, in the first and third scenario in the first row the difference between the distributions is very local), since this makes the comparison between the two aggregation methods interesting. We considered symmetric as well as asymmetric distributions. Gaussian shift and scale setups were considered in Appendix G. Such setups are less interesting in the context of this work, because if the two distributions differ only in shift or scale then specialized tests such as Wilcoxon rank-sum for shift will be preferable, but it is important to know that the suggested tests do not break down in this case. We used 20,000 simulated data sets, in each of the configurations of Figure 2.

Table 1 and Figure 3 show the power for the setups in Figure 2. These results show that if the number of intersections of the two densities is at least four, tests statistics with $m \geq 4$ have an advantage. Since the classical competitors, KS, CVM and AD, are based on $m = 2$, they perform far worse in these setups. Moreover, although HHG and DS have better power than the classical tests, HHG is essentially an $m \leq 3$ test, and DS penalizes large m s severely, therefore their power is still too low when fine partitioning is advantageous. The minimum p -value statistic, which does not require to preset m , is remarkably efficient: in Figure 2 we see that in all settings considered, it is close to the power of the optimal m .

The choice of aggregation by maximization versus summation depends on how local the differences are between the distributions. In Figure 3 we see clearly that when the differences are in very local areas, maximization achieves the greatest power and the test based on minimum p -value has more power if the aggregation is by maximization rather than by summation (setups 1, 2, and 13), and aggregation by summation is better otherwise. Note that the optimal m for aggregation by summation is always larger than for aggregation by maximization. The reason is that in order to have a powerful statistic aggregated by maximization, it is enough to have one good partition (i.e., contain cells where the distributions clearly differ) for a fixed m , whereas by summation it is necessary to have a large fraction of good partitions among all partitions of size m .

4.2 The Independence Problem

We examined the power properties of the ADP and DDP statistics, aggregated by summation for $m \in \{1, \dots, \sqrt{N}\}$, aggregated by maximization for $m \leq 4$, as well as the minimum p -value statistic $\min_{\{m \in 2, \dots, m_{\max}\}} p_m$ based on aggregation by summation. We display here the results for $m_{\max} = 10$ and $N = 100$.

We compared these tests to seven tests of independence suggested in the literature. We compared to Spearman's ρ , since it is perhaps the most widely used test to detect monotone associations. We also compared to previous tests suggested in the literature with the same two important properties as our suggested tests, namely proven consistency and distribution-freeness, as well as an available implementation: the test of Hoeffding (1948b), referred to as Hoeffding; the tests of Székely et al. (2007) and Heller et al. (2013) that first transform the observations of each variable into ranks, referred to as dCov and HHG, respectively. We also compared to the test of Reshef et al. (2011),

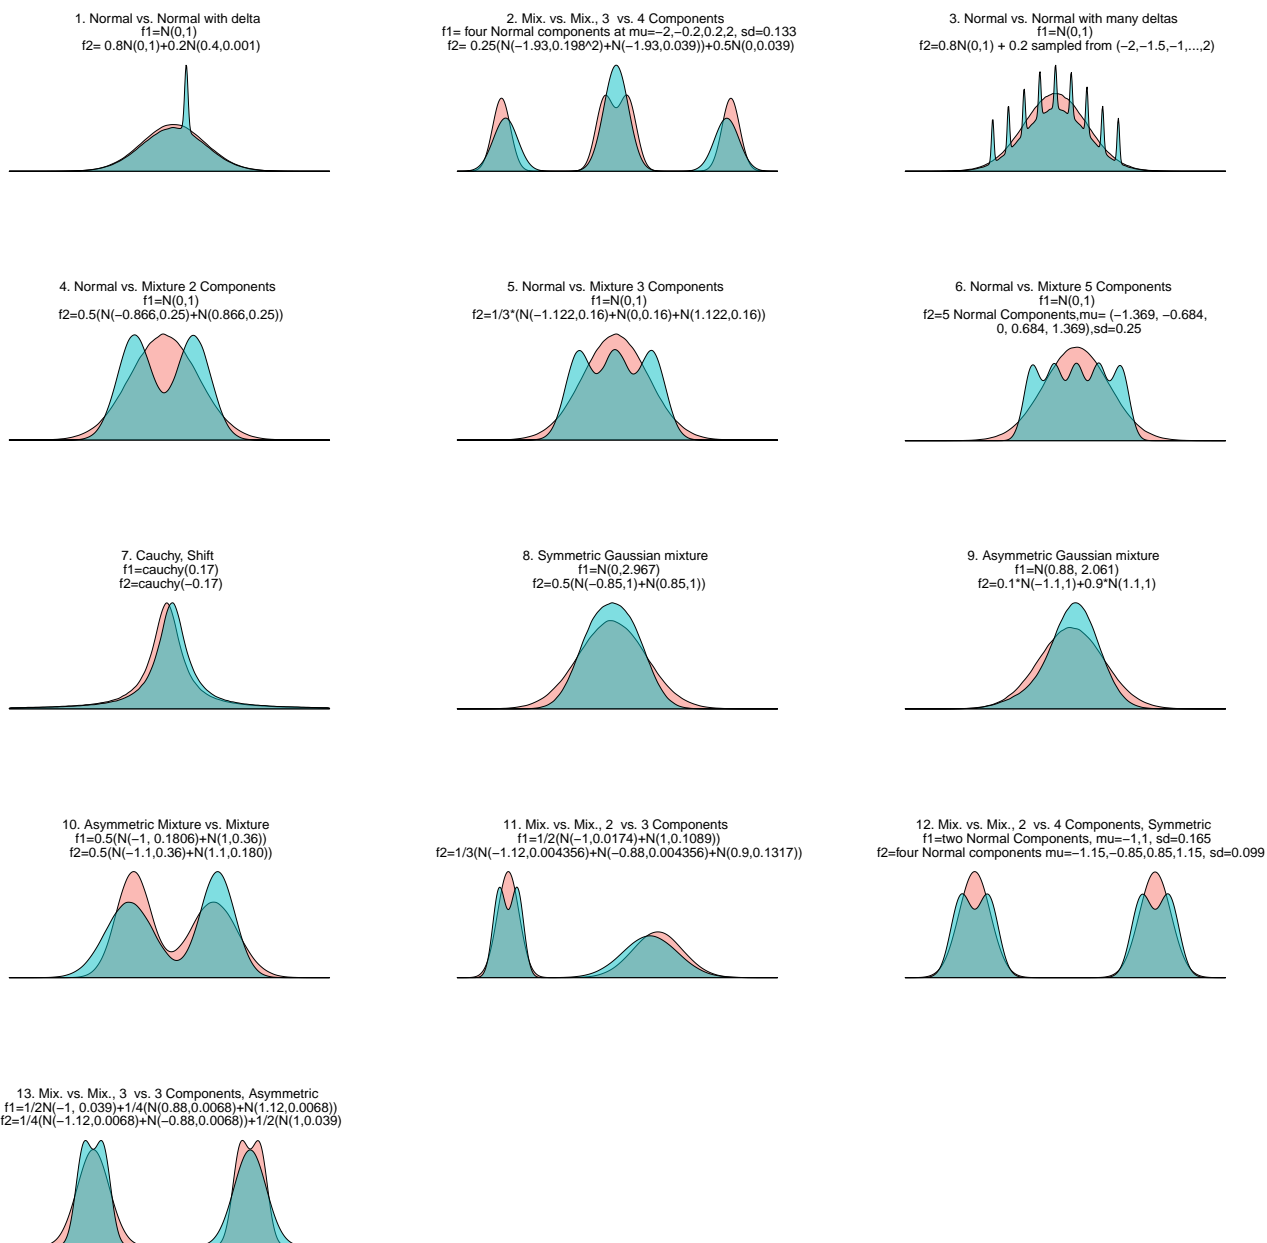


Figure 2: The two-sample problem in 13 different setups considered for $N = 500$, which differ in the number of intersections of the densities, the range of support where the differences lie, and whether they are symmetric or not.

DISTRIBUTION-FREE K -SAMPLE AND INDEPENDENCE TESTS

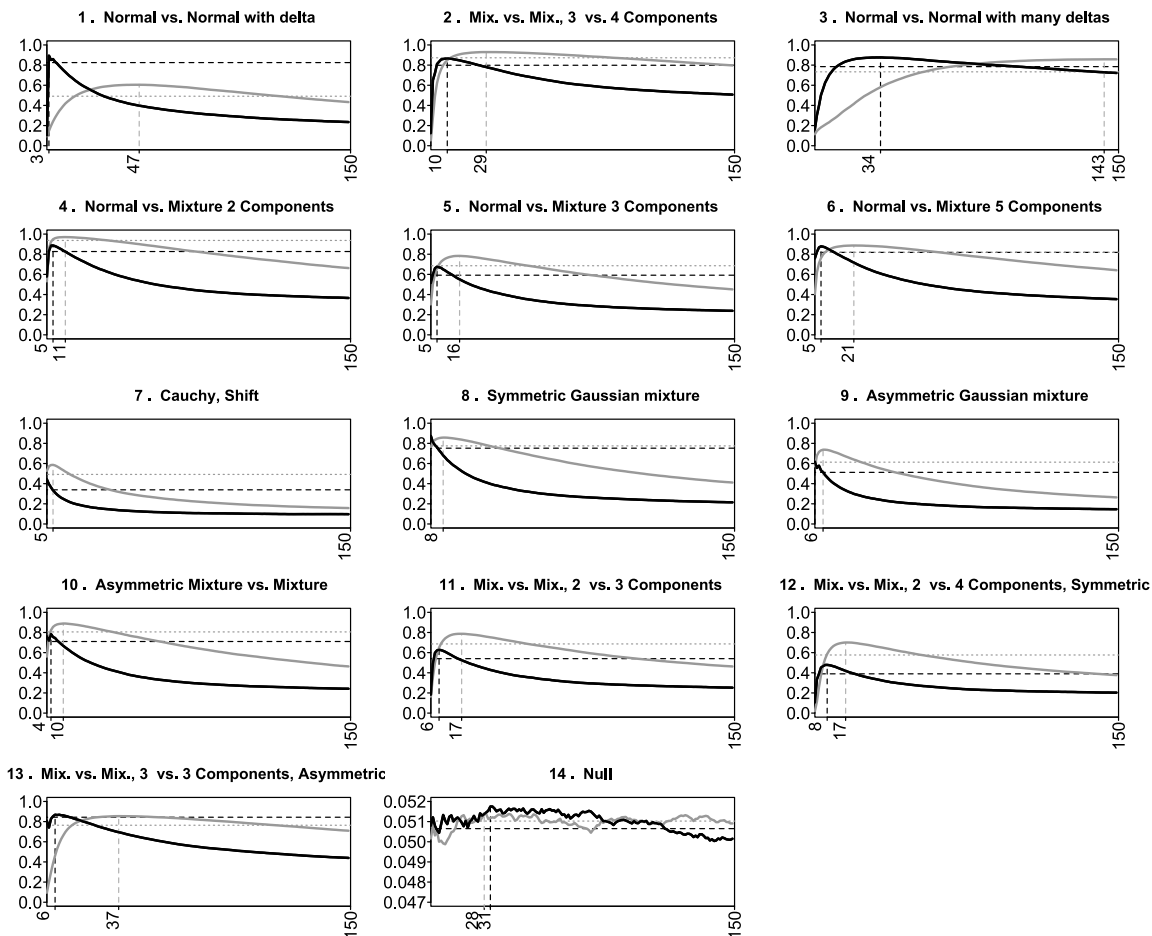


Figure 3: Estimated power with $N = 500$ sample points for the M_m (black) and S_m (grey) statistics for $m \in \{2, \dots, 149\}$ for the setups of Figure 2. The score per partition was the likelihood ratio test statistic. The power of the minimum p -value is the horizontal dashed black line when it combines the p -values based on M_m , and the horizontal dotted grey line when it combines the p -values based on S_m . The vertical lines show the optimal m for M_m (black) and S_m (grey).

Setup	Min. p -value aggregation		Wilcox	KS	CVM	AD	HHG	DS	
	by Max	by Sum							
1	Normal vs. Normal with delta	0.825	0.491	0.072	0.149	0.108	0.099	0.175	<u>0.849</u>
2	Mix. Vs. Mix., 3 Vs. 4 Components	0.799	<u>0.873</u>	0.000	0.020	0.001	0.021	0.344	0.560
3	Normal vs. Normal with many deltas	<u>0.785</u>	0.733	0.051	0.078	0.073	0.099	0.142	0.245
4	Normal vs. Mixture 2 Components	0.827	<u>0.937</u>	0.053	0.531	0.458	0.495	0.855	0.796
5	Normal vs. Mixture 5 Components	0.592	<u>0.686</u>	0.048	0.238	0.179	0.246	0.484	0.556
6	Normal vs. Mixture 10 Components	0.818	<u>0.820</u>	0.048	0.240	0.211	0.310	0.561	0.789
7	Cauchy, Shift	0.339	0.492	0.542	0.620	0.627	0.577	<u>0.641</u>	0.436
8	Symmetric Gaussian mixture	0.752	0.775	0.033	0.194	0.242	0.617	0.749	<u>0.835</u>
9	Asymmetric Gaussian mixture	0.512	0.613	0.050	0.253	0.277	0.469	<u>0.678</u>	0.599
10	Asymmetric Mixture vs. Mixture	0.711	<u>0.806</u>	0.000	0.159	0.119	0.395	0.690	0.747
11	Mix. Vs. Mix., 2 Vs. 3 Components	0.540	<u>0.686</u>	0.004	0.093	0.057	0.116	0.302	0.440
12	Mix. Vs. Mix., 2 Vs. 4 Comp., Sym.	0.390	<u>0.577</u>	0.000	0.005	0.000	0.005	0.079	0.270
13	Mix. Vs. Mix., 3 Vs. 3 Comp., Asym.	<u>0.844</u>	0.764	0.000	0.001	0.000	0.013	0.042	0.780
14	Null	0.051	0.051	0.050	0.043	0.050	0.050	0.050	0.050

Table 1: Power of competitors (columns 4-9), along with the minimum p -value statistic using the M_m p -values (column 2) and the S_m p -values (column 3), for $N = 500$. The score per partition was the likelihood ratio test statistic. The standard error was at most 0.0035. The advantage of the test based on the minimum p -value is large when the number of intersections of the two densities is at least four (setups 2,3,4,5,6,10,11,12, and 13). The best competitors are HHG and DS, but HHG is essentially an $m \leq 3$ test, and DS penalizes large m s severely, therefore in setups where $m \geq 4$ partitions are better they can perform poorly. Among the two variants in columns 2 and 3, the better choice clearly depends on the range of support in which the differences in distributions occur: aggregation by maximum has better power when the difference between the distributions is very local (setups 1, 3, and 13), and aggregation by summation has better power otherwise. The highest power per row is underlined.

referred to as MIC, which is not consistent due to the computational shortcuts they have to use. We note that the power of the original dCov and HHG was fairly similar to the power of their distribution-free variants, see Appendix I.

We examine complex bivariate relationships depicted in Figure 4. Most of these scenarios were collected from the literature illustrating the performance of other methods. Specifically, the first two rows were examined in Newton (2009), the next two rows are similar to the relationships examined in Reshef et al. (2011), and the Heavisine and Doppler examples in the last row were used extensively in the literature on denoising, see e.g., Donoho and Johnstone (1995). In all but the 4 Independent Clouds setup, there is dependence. The 4 Independent Clouds setup allows us to verify that the tests maintain the nominal level. We used 2000 simulated data sets for $N = 100$ and $N = 300$, in each of the configurations of Figure 4. Monotone setups are presented in Appendix H in Figure 12. Monotone setups are less interesting in the context of this work, because if there is reason to believe that the dependence is monotone, specialized tests such as Spearman’s ρ or Kendall’s τ will be preferable, but it is important to know that the suggested tests have reasonable power, as demonstrated in the results in Appendix H, Figure 13 and Table 10.

Tables 2 and 3, and Figure 5 show the power for the settings depicted in Figure 4. We only considered the test based on the DDP minimum p -value statistic in Tables 2 and 3, since for the

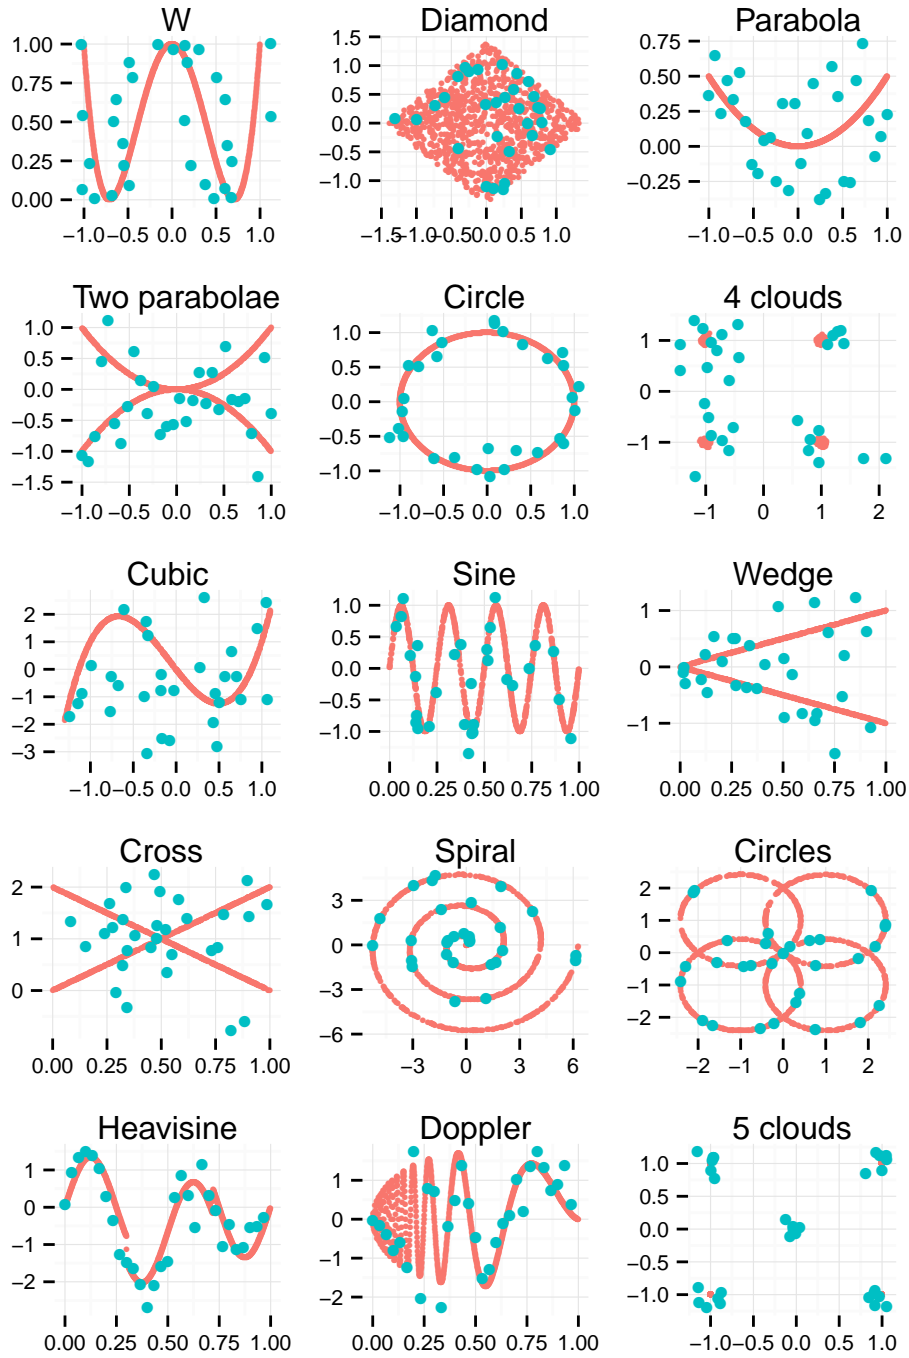


Figure 4: Bivariate relationships (in red), along with a sample $N = 100$ noisy observations (in blue). The “four clouds” relationship is a null relationship, where the two random variables are independent.

Setup	$\min_{\{m \in \{2, \dots, 10\}\}} p_m$	Spearman	Hoeffding	MIC	dCov	HHG
W	0.655	0.000	0.414	0.526	0.361	<u>0.798</u>
Diamond	0.919	0.013	0.116	0.074	0.074	<u>0.965</u>
Parabola	<u>0.847</u>	0.028	0.413	0.211	0.386	0.784
2Parabolas	<u>0.844</u>	0.095	0.135	0.048	0.124	0.723
Circle	<u>0.886</u>	0.000	0.033	0.046	0.002	0.850
Cubic	0.731	0.344	0.655	0.515	0.627	<u>0.768</u>
Sine	0.995	0.368	0.494	<u>1.000</u>	0.415	0.804
Wedge	0.595	0.064	0.360	0.303	0.338	<u>0.673</u>
Cross	0.704	0.089	0.160	0.069	0.130	<u>0.706</u>
Spiral	<u>0.949</u>	0.112	0.141	0.251	0.140	<u>0.337</u>
Circles	<u>0.999</u>	0.048	0.084	0.093	0.061	0.354
Heavisine	<u>0.710</u>	0.396	0.493	0.532	0.492	0.585
Doppler	0.949	0.513	0.784	<u>0.975</u>	0.744	0.912
5Clouds	<u>0.996</u>	0.000	0.000	0.561	0.004	0.904
4Clouds	0.051	0.050	0.057	0.050	0.051	0.050

Table 2: Power of competitors (columns 3–7), along with the DDP minimum p -value statistic for $N = 100$. The standard error is at most 0.011. The score per partition was the likelihood ratio test statistic. The DDP minimum p -value statistic performs very well in comparison to the other tests. Although the competitors may have greater power in some examples, the advantage is usually small. By far the best competitor is HHG, yet it has a disadvantage when the relationship is more complex, thus benefiting from the finer partition of the minimum p -value test, especially in the Sine, Heavisine, Spiral and Circles examples. The highest power per row is underlined.

minimum p -value statistic the tests of ADP and DDP are almost identical. These results provide strong evidence that for non-monotone noisy dependencies our tests have excellent power properties. Specifically, $S_{m \times m}^{DDP}$ with $m \in \{3, \dots, 10\}$ is more powerful than all other tests in Table 2 in most settings. For example, it had greater power than all competitors in 9 settings with $m = 4$ and in 11 settings with $m = 5$, out of the 14 non-null settings. The test based on the minimum p -value has greater power than all competitors in 7 settings, and it is very close to the best competitor in most of the other settings. The MIC is best for the Sine example but performs poorly in all other examples. The minimum p -value statistic is a close second best in the Sine example, with a difference of only 0.005 from MIC, yet all other tests are more than 0.19 below MIC in power. Overall, the HHG test is the best competitor, but its power is lower than that of the minimum p -value statistic when the optimal m is greater than 4. Table 3 shows that when aggregating by maximization, the choice of m matters and the power is higher for $m > 2$. However, the minimum p -value statistic, which is aggregated by summation and considers finer partitions, is more powerful for most settings, and is a close second in the remaining settings.

5. Application to Real Data

We examine the co-dependence between pairs of genes on chromosome 1 in the yeast gene expression data set from Hughes et al. (2000), which contained $N = 300$ expression levels. After removing genes with missing values, we had 94 genes and a family of $\binom{94}{2} = 4371$ pairs to examine simultaneously. Each pair was tested for independence by the tests of Spearman, Hoeffding,

DISTRIBUTION-FREE K -SAMPLE AND INDEPENDENCE TESTS

Setup	$\min_{\{m \in \{2, \dots, 10\}\}} p_m$	$M_{2 \times 2}^{DDP}$	$M_{3 \times 3}^{DDP}$	$M_{4 \times 4}^{DDP}$	$M_{2 \times 2}^{ADP}$
W	<u>0.655</u>	0.190	0.637	0.574	0.155
Diamond	0.919	0.272	<u>0.931</u>	0.912	0.247
Parabola	0.847	0.533	<u>0.855</u>	0.803	0.597
2Parabolas	0.844	0.466	<u>0.907</u>	0.897	0.578
Circle	<u>0.886</u>	0.222	0.880	0.884	0.170
Cubic	<u>0.731</u>	0.496	0.654	0.653	0.496
Sine	0.995	0.768	0.958	<u>0.998</u>	0.774
Wedge	<u>0.595</u>	0.410	0.536	0.478	0.455
Cross	<u>0.704</u>	0.268	0.680	0.673	0.341
Spiral	<u>0.949</u>	0.116	0.489	0.764	0.189
Circles	<u>0.999</u>	0.085	0.606	0.844	0.088
Heavisine	<u>0.710</u>	0.519	0.642	0.692	0.534
Doppler	0.949	0.828	0.969	<u>0.977</u>	0.833
5Clouds	0.996	0.062	<u>0.999</u>	<u>0.999</u>	0.076
4Clouds	0.051	0.052	0.051	0.051	0.052

Table 3: The power of different variants aggregated by maximization (columns 3–6), along with the DDP minimum p -value statistic (column 2) for $N = 100$. The standard error is at most 0.011. The score per partition was the likelihood ratio test statistic. Although maximization is better than summation in some examples, the advantage is usually small. The advantage of the minimum p -value statistic, which is based on aggregation by summation, is large in the Cubic, Cross, Spiral, and Circles relationships. In most examples, power increases with m . The power differences between the ADP and DDP variants are small. We present only the maximum variants that take at most $O(N^3)$ to compute, therefore for ADP only results with $m = 2$ are presented. The highest power per row is underlined.

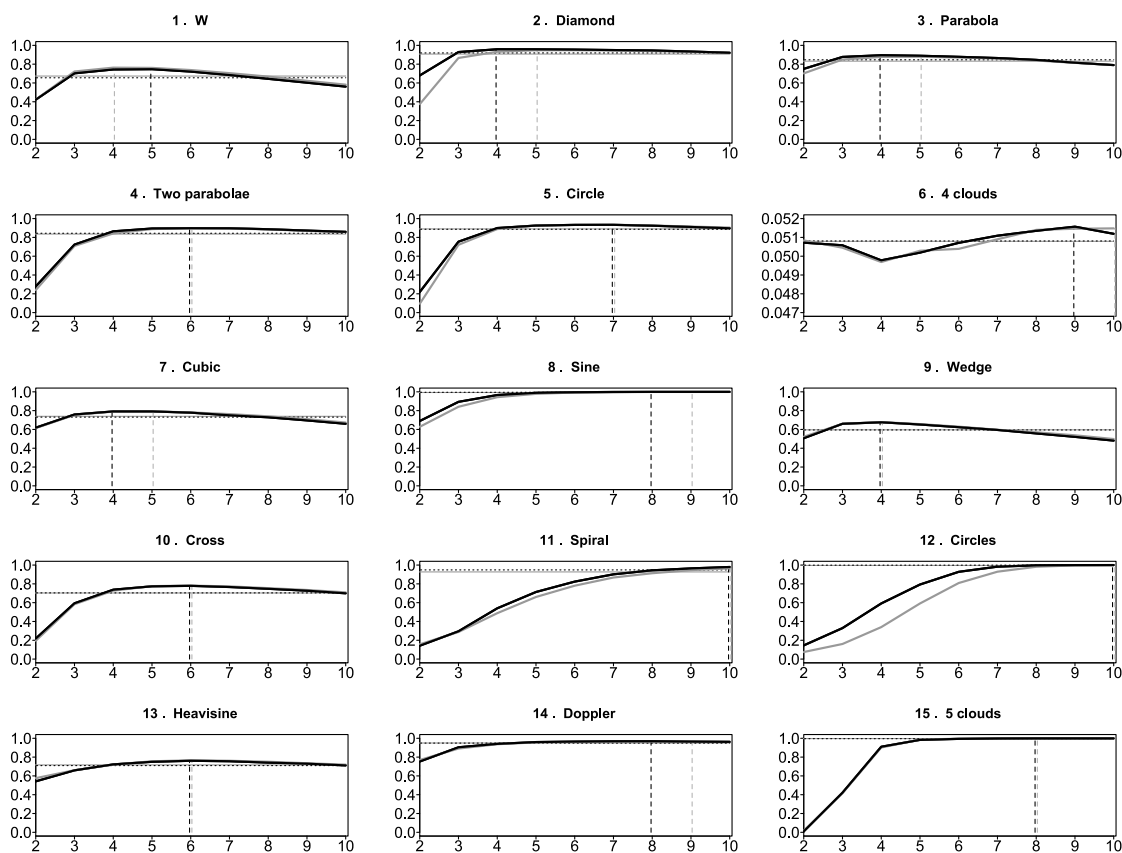


Figure 5: Estimated power as a function of partition size m , with $N = 100$ sample points, for the DDP (black) and ADP (grey) summation variants using the likelihood ratio score for the setups of Figure 4. The score per partition was the likelihood ratio test statistic. For DDP (black) and ADP (grey), the horizontal dashed line is the power of the minimum p -value statistic, and the vertical lines is the optimal m .

Test	Number of rejections	Number of intersections
Spearman	2488	2445
MIC	245	245
Hoeffding	2890	2844
HHG on ranks	3283	3199
dCov on ranks	2889	2845
minimum p -value based on ADP	3310	3294

Table 4: Benjamini–Hochberg rejections at level 0.05 for the gene expression problem of Hughes et al. (2000). For different test statistics (rows), the number of rejections (column 2), and their intersection with the rejections using the minimum p -value statistic on DDP (column 3). The minimum p -value statistic on DDP had the highest number of rejections, 3312.

MIC, dCov and HHG on ranks, as well as by our new tests with m ranging from 2 to $m_{\max} = 17$. The null tables were based on 20,000 permutations for $N = 300$. The adjusted p -values from the Benjamini–Hochberg procedure (Benjamini and Hochberg, 1995) were computed for each test statistic.

Table 4 shows the pairwise agreements between the Benjamini–Hochberg procedure at level 0.05 using the different test statistics considered in each row, with the minimum p -value statistic based on DDP. Clearly, a large number of pairwise associations are missed when testing is performed with Spearman’s ρ compared to the minimum p -value statistic, and only a small number of gene pairs detected with Spearman are missed by the minimum p -value statistic (row 1 in Table 4). These findings contradict an earlier examination of the data. Steuer et al. (2002) concluded that the most widely used approach for pairwise association testing, namely Spearman correlation, performs equivalently to a mutual information based testing approach. The authors speculated that actual dependencies, if any, are linear. The number of discoveries using MIC, Hoeffding, and dCov are much smaller than using the minimum p -value statistic. HHG on ranks also discovers less co-dependencies compared with the minimum p -value statistic. The agreement between the tests based on DDP and ADP was very high, as seen in the last row of Table 4 and in Figure 6, which shows the number of rejections for $S_{m \times m}^{DDP}$ and $S_{m \times m}^{ADP}$ for $m = 2, \dots, 17$. We conclude that in this data set there are many nonlinear associations, but powerful tests are necessary in order to detect such associations in light of the large number of simultaneous tests that have to be carried out, and that the suggested tests can be valuable tools for this task.

Note that the data had ties due to low precision of the documented expression levels. Ties were broken randomly, see remark 2.1. Repeated analysis with different seeds provided similar results.

6. Discussion

In this paper we proposed new partition-based test statistics for both the independence problem and the two-sample problem. We proved that the statistics are consistent for general alternatives and demonstrated in simulations that the power advantage of the tests based on finer partitions can be great. We further showed that the power of our regularized statistics is very close to that of the statistics based on the optimal partition size. We recommend the test using the minimum p -value statistic based on aggregation by summation, unless the alternative is suspected to be of very local nature. Specifically, in the K -sample problem if the difference between the distributions is on a




Figure 6: Number of discoveries of genes with associated expression patterns in the data of Hughes et al. (2000), by the Benjamini–Hochberg procedure at level 0.05 using $S_{m \times m}^{DDP}$ (black) and $S_{m \times m}^{ADP}$ (gray) for $m = 2, \dots, 17$. In addition, number of rejections using the minimum p -value statistic with $m_{\max} = 17$, using the DDP (black horizontal line) and the ADP (gray horizontal line).

very small range of the support, then the aggregation by maximization is preferred over aggregation by summation.

The algorithms described in Section 3.1 for the test of independence based on regularized scores for a range of $m \times m$ partitions can easily be generalized to include $m_x \times m_y$ partitions, where $m_x \neq m_y$, with the same complexity for the ADP statistic (for the DDP statistic $m_x = m_y$). Considering unequal partition sizes for X and Y is expected to improve power when the (unknown) optimal partition has an m_x value very different than the m_y value. Moreover, when the (unknown) optimal partition has $m_x \approx m_y$, the power loss from considering the minimum p -value over all $m_x \times m_y$ values instead of over all $m \times m$ values is expected to be small.

The algorithms we suggested for the K -sample problem are $O(N^2)$ and therefore are feasible even for large N . For the test of independence, even though the complexity of our suggested algorithms is $O(N^4)$, for small N these algorithms can be quite efficient in the following quite common multiple testing setting in modern studies. If M hypotheses are simultaneously examined with the same sample size, then the computational complexity of using our distribution-free tests is $O(M \times N^4)$ if null table is available for this N , or $\max\{O(M \times N^4), O(B \times N^4)\}$ if the null table is generated by the user using B Monte-Carlo replicates for the sample size N . If we needed to recompute the null distribution for every one of the M hypotheses (as required for permutation tests that are not distribution-free, such as dCov and HHG), then the computational complexity would have been $O(M \times B \times N^4)$, which may be infeasible in modern studies where the number of hypotheses tests simultaneously examined can be several thousands or hundreds of thousands. Since the null distribution needs to be generated only once in order to compute the significance

of all the M test statistics, due to the distribution-free property of our tests, they can be feasible with today's computing power even for a few thousands samples. However, computing $O(N^4)$ test statistics is unfeasible for larger sample sizes. To reduce the computational complexity when N is large, statistics which do not go over all partitions but rather just over a representative sample can be considered. This approach was used for example in Jiang (2014) for the K -sample problem. A simple way of doing this is to divide the data into $\sqrt{N} \times \sqrt{N}$ bins and only consider partitions that do not break up these bins. We expect such statistics to also be consistent and the algorithms that accompany them to be computable in $O(N^2)$.

If one expects relatively simple dependence structures, for large N , the $S_{3 \times 3}^{DDP}$ is recommended, since it is both distribution-free and computable in $O(N^2)$ (see Remark 3.1). In our simulations it was as powerful as HHG and more powerful than dCov, and it has the advantage of being distribution-free.

We focused our attention mainly on comparisons of distribution-free tests. We provided in Appendix I a comparison of our recommended tests with the non distribution-free HHG test of Heller et al. (2013) and the dCov test of Székely and Rizzo (2009b). These tests had less power than our novel distribution-free tests when there is advantage to finer partitions. It is of interest to examine how the extended α -distance dependence measure suggested in Székely and Rizzo (2009b) perform in these complex relationships. It is also of interest to compare to the non distribution-free kernel methods of Gretton et al. (2007) and Gretton et al. (2008). For this purpose, care must be taken in the choice of kernel bandwidth, see Gretton et al. (2012b).

In this work we restricted our attention to testing an independent sample from the joint distribution. If the paired observations are not independent, but are independent within blocks, we can easily modify our test by restricting the permutation tests to be within the blocks. When there is temporal dependence, an interesting work by Chwialkowski et al. (2014) suggested a bootstrap based approach for kernel based tests. How to modify our test in the face of temporal dependence is an open question.

A thorough investigation of the suggested mutual information estimator in Section 2 was outside the scope of this paper, but is of interest for future research. We suspect the asymptotic distribution of our mutual information estimator has a simple form. The bias of the estimator can be dealt with by modifying our estimator, and our algorithms accordingly, to only include partitions with cells of a minimum size, and by bias correction methods suggested in the literature, e.g., Vu et al. (2007). Although in this work we limited ourselves to a theoretical examination of the ADP summation statistic for mutual information estimation, we recognize that an estimator based only on the DDP may be useful, and we plan to explore it in the future.

Acknowledgments

R. Heller and B. Brill are supported by Israel Science Foundation grant 2012896. S. Kaufman is supported by a fellowship from the Edmond J. Safra Center for Bioinformatics at Tel Aviv University.

Appendix A. Proof of Theorem 2

The proof for the DDP and ADP tests are given in Sections A.1 and A.2 respectively.

A.1 The DDP Test

Proof Denote $S_{m \times m} = S_{m \times m}^{DDP}$. For simplicity, we show the proof using Pearson's test statistic. The proof using the likelihood ratio test statistic is very similar and therefore omitted. We want to show that for an arbitrary fixed $\alpha \in (0, 1)$, if H_0 is false, then $\lim_{N \rightarrow \infty} Pr(S_{m \times m} > S_{1-\alpha}^{tab}) = 1$, where $S_{1-\alpha}^{tab}$ denotes the $1 - \alpha$ quantile of the null distribution of $S_{m \times m}$.

If H_0 is false, then without loss of generality $h(x_0, y_0) > f(x_0)g(y_0)$. Moreover, there exists a distance $R > 0$ such that $h(x, y) > f(x)g(y)$ for all points (x, y) in the set $\mathcal{A} = \{(x, y) : x_0 \leq x \leq x_0 + R, y_0 \leq y \leq y_0 + R\}$. The set \mathcal{A} has positive probability, and moreover

$$\min_{\mathcal{A}} [h(x, y) - f(x)g(y)] > 0.$$

Denote this minimum by $c > 0$. Clearly, the following two subsets of \mathcal{A} have positive probability as well:

$$\begin{aligned} \mathcal{A}_1 &= \{(x, y) : x_0 \leq x \leq x_0 + R/4, y_0 \leq y \leq y_0 + R/4\} \\ \mathcal{A}_2 &= \{(x, y) : x_0 + 3R/4 \leq x \leq x_0 + R, y_0 + 3R/4 \leq y \leq y_0 + R\}. \end{aligned}$$

Denote the probabilities of \mathcal{A}_1 and \mathcal{A}_2 by f_1 and f_2 , respectively.

Let $\Gamma\{(x_1, y_1), \dots, (x_N, y_N)\}$ be the set of partitions of size m based on at least one sample point in \mathcal{A}_1 and on at least one sample point in \mathcal{A}_2 . Let N_i denote the number of sample points in \mathcal{A}_i , $i \in \{1, 2\}$. Let $\mathcal{I} \in \Gamma\{(x_1, y_1), \dots, (x_N, y_N)\}$ be such a (arbitrary fixed) partition. So for \mathcal{I} there exists $(i, j) \subseteq \mathcal{I}$ such that $(x_i, y_i) \in \mathcal{A}_1$ and $(x_j, y_j) \in \mathcal{A}_2$. Consider the cell C defined by the two points (i, j) .

The fraction of observed counts in the cell C is a linear combination of empirical cumulative distribution functions

$$\frac{o_C}{N - (m - 1)} = \hat{F}_{XY}(x_i, y_i) + \hat{F}_{XY}(x_j, y_j) - \hat{F}_{XY}(x_i, y_j) - \hat{F}_{XY}(x_j, y_i),$$

and the expected fraction under the null is a function of the marginal cumulative distributions

$$\frac{e_C}{N - (m - 1)} = \{\hat{F}_X(x_j) - \hat{F}_X(x_i)\}\{\hat{F}_Y(y_j) - \hat{F}_Y(y_i)\}.$$

where \hat{F} denotes the empirical distribution function based on $N - (m - 1)$ sample points.

By the Glivenko-Cantelli theorem, uniformly almost surely,

$$\begin{aligned} \lim_{N \rightarrow \infty} \left(\frac{o_C}{N - (m - 1)} - \int_{\{(x, y) : x \in (x_i, x_j], y \in (y_i, y_j]\}} h(x, y) dx dy \right) &= 0, \\ \lim_{N \rightarrow \infty} \left\{ \frac{e_C}{N - (m - 1)} - \left(\int_{\{x : x \in (x_i, x_j]\}} f(x) dx \right) \left(\int_{\{y : y \in (y_i, y_j]\}} g(y) dy \right) \right\} &= 0. \end{aligned} \tag{13}$$

It is important that the convergence be uniform in order to obtain a positive lower bound on the average over all partitions, see the discussion before equation (18).

Therefore, by Slutsky's theorem and the continuous mapping theorem, we have that uniformly almost surely

$$\begin{aligned}
 \lim_{N \rightarrow \infty} \frac{1}{N - (m - 1)} \frac{(o_C - e_C)^2}{e_C} &= \lim_{N \rightarrow \infty} \frac{\left(\frac{o_C}{N - (m - 1)} - \frac{e_C}{N - (m - 1)} \right)^2}{\frac{e_C}{N - (m - 1)}} \\
 &\geq \lim_{N \rightarrow \infty} \left(\frac{o_C}{N - (m - 1)} - \frac{e_C}{N - (m - 1)} \right)^2 \\
 &= \lim_{N \rightarrow \infty} \left[\int_{\{(x,y):x \in (x_i, x_j], y \in (y_i, y_j]\}} \{h(x, y) - f(x)g(y)\} dx dy \right]^2, \tag{14}
 \end{aligned}$$

where the inequality follows from the fact that $\frac{e_C}{N - (m - 1)} \leq 1$.

We shall show that this limit can be bounded from below by a positive constant that depends on (x_0, y_0) and R but not on \mathcal{I} . Since

$$\begin{aligned}
 &\{(x, y) : x \in (x_0 + R/4, x_0 + 3R/4], y \in (y_0 + R/4, y_0 + 3R/4]\} \\
 &\subseteq \{(x, y) : x \in (x_i, x_j], y \in (y_i, y_j]\},
 \end{aligned}$$

a positive lower bound on expression (14) can be obtained:

$$\begin{aligned}
 &\lim_{N \rightarrow \infty} \left[\int_{\{(x,y):x \in (x_i, x_j], y \in (y_i, y_j]\}} \{h(x, y) - f(x)g(y)\} dx dy \right]^2 \\
 &\geq \left[\int_{\{(x,y):x \in (x_0 + R/4, x_0 + 3R/4], y \in (y_0 + R/4, y_0 + 3R/4]\}} \{h(x, y) - f(x)g(y)\} dx dy \right]^2 \\
 &\geq c^2 \int_{\{(x,y):x \in (x_0 + R/4, x_0 + 3R/4], y \in (y_0 + R/4, y_0 + 3R/4]\}} dx dy = c^2 R^2 / 4, \tag{15}
 \end{aligned}$$

where the first inequality follows since $h(x, y) - f(x)g(y) > 0$ in \mathcal{A} , and the second inequality follows since the minimum value is $c > 0$. Therefore, it follows that $\frac{1}{N - m + 1} \frac{(o_C - e_C)^2}{e_C}$ converges uniformly almost surely to a positive constant greater than $c' = c^2 R^2 / 4$,

$$Pr \left(\lim_{N \rightarrow \infty} \frac{1}{N - m + 1} \frac{(o_C - e_C)^2}{e_C} \geq c' \right) = 1. \tag{16}$$

The partition \mathcal{I} either contains the cell C , or a group of cells that divide C . By Jensen's inequality, it follows that if the partition \mathcal{I} contains a group of cells that divide C , the score is made larger, since for any partition of the cell C , $C = \cup_l C_l$,

$$\left(\frac{o_C - e_C}{e_C} \right)^2 = \left(\frac{\sum_l e_{C_l} \left(\frac{o_{C_l}}{e_{C_l}} - 1 \right)}{\sum_h e_{C_h}} \right)^2 \leq \frac{\sum_l e_{C_l} \left(\frac{o_{C_l}}{e_{C_l}} - 1 \right)^2}{\sum_l e_{C_l}} = \frac{\sum_l \frac{(o_{C_l} - e_{C_l})^2}{e_{C_l}}}{e_C},$$

and therefore

$$\frac{(o_C - e_C)^2}{e_C} \leq \sum_l \frac{(o_{C_l} - e_{C_l})^2}{e_{C_l}}. \quad (17)$$

Since $\sum_l \frac{(o_{C_l} - e_{C_l})^2}{e_{C_l}}$ or $\frac{(o_C - e_C)^2}{e_C}$ is part of the sum that defines $T^{\mathcal{I}}$, it follows from equations (17) and (16) that $\frac{T^{\mathcal{I}}}{N-m+1}$ converges uniformly almost surely to a positive constant greater than c' .

Let $|\Gamma|$ denote the cardinality of $\Gamma\{(x_1, y_1), \dots, (x_N, y_N)\}$. Since $\mathcal{I} \in \Gamma\{(x_1, y_1), \dots, (x_N, y_N)\}$ was arbitrary fixed, and since the convergence for fixed \mathcal{I} of $t_C/[N - (m - 1)]$ to a limit bounded from below by a positive constant was uniform, it follows that $\frac{1}{|\Gamma|} \sum_{\mathcal{I} \in \Gamma} \frac{T^{\mathcal{I}}}{N-m+1}$ converges almost surely to a positive constant at least as large as c' . To see this, note that from the uniform convergence in equation (16), it follows that for an arbitrary fixed $\epsilon > 0$, there exists $N(\epsilon)$ (which does not depend on C) such that for all $N > N(\epsilon)$, $\frac{1}{N-(m-1)} \frac{(o_C - e_C)^2}{e_C} \geq c' - \epsilon$ for every C , and therefore that $\frac{1}{|\Gamma|} \sum_{\mathcal{I} \in \Gamma} \frac{T^{\mathcal{I}}}{N-m+1} \geq c' - \epsilon$ for all $N > N(\epsilon)$.

Since $S_{m \times m} \geq \sum_{\mathcal{I} \in \Gamma} T^{\mathcal{I}}$, it follows that almost surely

$$\lim_{N \rightarrow \infty} \frac{S_{m \times m}}{|\Gamma|(N - m + 1)} > c'. \quad (18)$$

We shall show that $\lim_{N \rightarrow \infty} |\Gamma|/\binom{N}{m-1}$ is bounded below by a positive constant. First, we shall consider the case where m is finite. Then, a subset of $\Gamma\{(x_1, y_1), \dots, (x_N, y_N)\}$ is the set of all partitions with $m - 2$ sample points in \mathcal{A}_1 , and one sample point in \mathcal{A}_2 . Therefore, $|\Gamma| \geq \binom{N_1}{m-2} N_2$.

Simple algebraic manipulations lead to the following expression for $\frac{\binom{N_1}{m-2} N_2}{\binom{N}{m-1}}$:

$$(m-1) \frac{N_2}{N} \frac{N_1}{N-1} \cdots \frac{N_1 - m + 3}{N - m + 2}.$$

Since N_1/N converges almost surely to f_1 , and similarly N_2/N converges almost surely to f_2 , then for $m \geq 3$ finite it follows that $\frac{\binom{N_1}{m-2} N_2}{\binom{N}{m-1}}$ converges almost surely to a positive constant. Therefore, $|\Gamma|/\binom{N-1}{m-1}$ is bounded away from zero. Second, we shall consider the case that $m \rightarrow \infty$. The complement of Γ , Γ^C , is the set of contingency tables with no points in \mathcal{A}_1 or in \mathcal{A}_2 . An upper bound for $|\Gamma^C|$ is

$$\binom{N - N_1}{m-1} + \binom{N - N_2}{m-1}.$$

Note that in order to show that $\lim_{N \rightarrow \infty} |\Gamma|/\binom{N}{m-1}$ is bounded below by a positive constant, since $|\Gamma| = \binom{N}{m-1} - |\Gamma^C|$, it is enough to show that $|\Gamma^C|/\binom{N}{m-1}$ converges to zero as $N \rightarrow \infty$. Simple algebraic manipulations lead to the following expression for $\frac{\binom{N-N_1}{m-1}}{\binom{N}{m-1}}$:

$$\left(1 - \frac{N_1}{N}\right) \cdots \left(1 - \frac{N_1}{N - (m-2)}\right) \leq \left(1 - \frac{N_1}{N}\right)^m.$$

Since N_1/N converges almost surely to $f_1 \in (0, 1)$, it follows that $\frac{\binom{N-N_1}{m-1}}{\binom{N}{m-1}}$ converges almost surely to zero as $m \rightarrow \infty$. Similarly, since N_2/N converges almost surely to $f_2 \in (0, 1)$, it follows that $\frac{\binom{N-N_2}{m-1}}{\binom{N}{m-1}}$ converges almost surely to zero. Therefore, $|\Gamma^C|/\binom{N}{m-1}$ converges to zero as $N \rightarrow \infty$.

Since $\lim_{N \rightarrow \infty} |\Gamma| / \binom{N}{m-1}$ is bounded below by a positive constant, it follows from (18) that almost surely

$$\lim_{N \rightarrow \infty} \frac{S_{m \times m}}{\binom{N}{m-1}(N-m+1)} \geq c'', \quad (19)$$

for some constant $c'' > 0$.

Consider now a random permutation $(\pi y_1), \dots, (\pi y_N)$ of the y -values y_1, \dots, y_N . Let $S_{m \times m}^\pi$ be the test statistic that is computed from the data $(x_1, \pi y_1), \dots, (x_N, \pi y_N)$. Therefore, by Markov's inequality,

$$\begin{aligned} Pr \left(S_{m \times m}^\pi \geq c'' \binom{N}{m-1} (N-m+1) \mid \vec{x}, \vec{y} \right) &\leq \frac{E(S_{m \times m}^\pi \mid \vec{x}, \vec{y})}{c'' \binom{N}{m-1} (N-m+1)} \\ &\approx \frac{\binom{N}{m-1} (m-1)^2}{c'' \binom{N}{m-1} (N-m+1)}. \end{aligned} \quad (20)$$

where $\vec{x} = (x_1, \dots, x_N)$ and $\vec{y} = (y_1, \dots, y_N)$. The approximation in (20) becomes more accurate the larger N is, since each of the contingency tables is approximately χ^2 with $(m-1)^2$ degrees of freedom. The right hand side of equation (20) goes to 0 as $N \rightarrow \infty$, as long as $\lim_{N \rightarrow \infty} \frac{m}{\sqrt{N}} = 0$. Thus,

$$\lim_{N \rightarrow \infty, m/\sqrt{N} \rightarrow 0} Pr \left(S_{m \times m}^\pi \geq c'' \binom{N}{m-1} (N-m+1) \mid \vec{x}, \vec{y} \right) = 0. \quad (21)$$

We now have all the necessary results to complete the proof. Specifically,

$$\lim_{N \rightarrow \infty} Pr(S_{m \times m} \leq S_{1-\alpha}^{tab}) \leq \lim_{N \rightarrow \infty} Pr \left\{ S_{m \times m} \leq c'' \binom{N}{m-1} (N-m+1) \right\} = 0,$$

where the inequality follows from (21), since $S_{1-\alpha}^{tab}$ is below $c'' \binom{N}{m-1} (N-m+1)$ for N large enough, and the equality follows from (19), thus proving item 1 of Theorem 2.

To prove item 2, we will use the following inequality for chi-square distributions, which appears in equation (4.3) of Laurent and Massart (2000): for U a χ^2 statistic with D degrees of freedom, for any positive x , $Pr(U - D \geq 2\sqrt{Dx} + 2x) \leq e^{-x}$.

Let \mathcal{I} be a fixed arbitrary partition of size m . Since for N large enough, under the null hypothesis, $T^\mathcal{I}$ is approximately a χ^2 statistic with $D = (m-1)^2$ degrees of freedom, it thus follows that for $x > D$,

$$Pr_{H_0}(T^\mathcal{I} - D \geq 4x) \leq e^{-x}. \quad (22)$$

Let $x = \frac{c'}{8}(N-m+1) - D/4$. Then for N large enough, $x > D$. It thus follows that

$$Pr_{H_0}(T^\mathcal{I} \geq \frac{c'}{2}(N-m+1)) \leq e^{-(\frac{c'}{8}(N-m+1) - D/4)} \leq e^{-(\frac{c'}{8}(N-m+1) - (m-1)^2/4)} \quad (23)$$

By Bonferroni's inequality,

$$\begin{aligned} Pr_{H_0} \left(M_{m \times m}^{DDP} \geq \frac{c'}{2}(N-m+1) \right) &\leq \sum_{\mathcal{I} \in \Pi_m^{DDP}} Pr_{H_0} \left(T^\mathcal{I} \geq \frac{c'}{2}(N-m+1) \right) \\ &\leq \binom{N}{m} e^{-(\frac{c'}{8}(N-m+1) - (m-1)^2/4)}, \end{aligned} \quad (24)$$

where the last inequality follows from (23). Since $\binom{N}{m}$ is at most $O(N^{\sqrt{N}})$, and since

$$e^{-(\frac{c'}{8}(N-m+1)-(m-1)^2/4)} = O(e^{-(\frac{c'}{8}N)}),$$

it follows that the expression in (24) goes to zero as $N \rightarrow \infty$.

Since we found contingency tables for which under the alternative the test statistic $\frac{T^{\mathcal{I}}}{N-m+1}$ converges uniformly almost surely to a positive constant greater than c' (16), it follows that $\frac{M_{m \times m}^{DDP}}{N-m+1}$ converges uniformly almost surely to a positive constant greater than c' when the null is false. From (24) it follows that as $N \rightarrow \infty$, with $\lim_{N \rightarrow \infty} \frac{m}{\sqrt{N}} = 0$, the probability that the test statistics $\frac{M_{m \times m}^{DDP}}{N-m+1}$ will be above $\frac{c'}{2}$ goes to zero when the null is true. It follows that the null hypothesis will be rejected with asymptotic probability one when it is false. \blacksquare

A.2 The ADP Test

Proof We want to show that if H_0 is false, then for an arbitrary fixed α , $\lim_{N \rightarrow \infty} Pr(S_{m \times m}^{ADP} > S_{1-\alpha}^{tab}) = 1$, where $S_{1-\alpha}^{tab}$ denotes the $1 - \alpha$ quantile of the null distribution of $S_{m \times m}^{ADP}$. We use $\mathcal{A}, c, \mathcal{A}_1, \mathcal{A}_2, f_1, f_2$ as defined in the beginning of Appendix A of the main text.

For the ADP test, recall that the partitioning is based on selecting $m-1$ points from $1.5, \dots, N-0.5$ for the partitions of the ranked x -values, and separately for the partitions of the ranked y -values. For a fixed rectangle, we say a grid point $(i+0.5, j+0.5)$ is in the rectangle if the two x -values with ranks i and $i+1$, and the two y -values with ranks j and $j+1$, are in the rectangle, for $(i, j) \in \{1, \dots, N\}^2$. Let $\Gamma\{(x_1, y_1), \dots, (x_N, y_N)\}$ be the set of partitions of size m with at least one grid point in \mathcal{A}_1 and at least one grid point in \mathcal{A}_2 . Let N_{ix} be the number of x -coordinates of the grid points in $\mathcal{A}_i, i \in \{1, 2\}$, and N_{iy} be the number of y -coordinates of the grid points in $\mathcal{A}_i, i \in \{1, 2\}$.

Let $\mathcal{I} \in \Gamma\{(x_1, y_1), \dots, (x_N, y_N)\}$ define an (arbitrary fixed) ADP partition in Γ . There exist two x -values in \mathcal{A}_1 that are separated by a grid point in \mathcal{I} , and two x -values in \mathcal{A}_2 that are separated by a grid point in \mathcal{I} , denote the average of these two x -values by x_1^* and x_2^* . Let y_1^* and y_2^* be similarly defined for the y -values.

Let C be the cell defined by the points $(x_i^*, y_i^*), i = 1, 2$. The fraction of observed counts in the cell C is a linear combination of empirical cumulative distribution functions

$$\frac{o_C}{N} = \hat{F}_{XY}(x_1^*, y_1^*) + \hat{F}_{XY}(x_2^*, y_2^*) - \hat{F}_{XY}(x_1^*, y_2^*) - \hat{F}_{XY}(x_2^*, y_1^*),$$

and the expected fraction under the null, is a function of the cumulative marginal distributions

$$\frac{e_C}{N} = \{\hat{F}_X(x_2^*) - \hat{F}_X(x_1^*)\}\{\hat{F}_Y(y_2^*) - \hat{F}_Y(y_1^*)\},$$

where \hat{F} denotes the empirical cumulative distribution function based on N sample points.

By the Glivenko–Cantelli theorem, uniformly almost surely,

$$\begin{aligned} \lim_{N \rightarrow \infty} \left(\frac{o_C}{N} - \int_{\{(x,y):x \in (x_1^*, x_2^*], y \in (y_1^*, y_2^*]\}} h(x, y) dx dy \right) &= 0, \\ \lim_{N \rightarrow \infty} \left\{ \frac{e_C}{N} - \left(\int_{\{x: x \in (x_1^*, x_2^*]\}} f(x) dx \right) \left(\int_{\{y: y \in (y_1^*, y_2^*]\}} g(y) dy \right) \right\} &= 0. \end{aligned} \quad (25)$$

Therefore, by Slutsky’s theorem and the continuous mapping theorem, we have that uniformly almost surely

$$\begin{aligned} \lim_{N \rightarrow \infty} \frac{1}{N} \frac{(o_C - e_C)^2}{e_C} &= \lim_{N \rightarrow \infty} \frac{\left(\frac{o_C}{N} - \frac{e_C}{N} \right)^2}{\frac{e_C}{N}} \\ &\geq \lim_{N \rightarrow \infty} \left(\frac{o_C}{N} - \frac{e_C}{N} \right)^2 \\ &= \lim_{N \rightarrow \infty} \left[\int_{\{(x,y):x \in (x_1^*, x_2^*], y \in (y_1^*, y_2^*]\}} \{h(x, y) - f(x)g(y)\} dx dy \right]^2, \end{aligned} \quad (26)$$

where the inequality follows from the fact that $\frac{e_C}{N} \leq 1$.

We shall show that the limit (26) can be bounded from below by a positive constant that depends on (x_0, y_0) and R but not on \mathcal{I} . Since

$$\begin{aligned} &\{(x, y) : x \in (x_0 + R/4, x_0 + 3R/4], y \in (y_0 + R/4, y_0 + 3R/4]\} \\ &\subseteq \{(x, y) : x \in (x_1^*, x_2^*], y \in (y_1^*, y_2^*]\}, \end{aligned}$$

a positive lower bound can be obtained:

$$\begin{aligned} &\lim_{N \rightarrow \infty} \left[\int_{\{(x,y):x \in (x_1^*, x_2^*], y \in (y_1^*, y_2^*]\}} \{h(x, y) - f(x)g(y)\} dx dy \right]^2 \\ &\geq \left[\int_{\{(x,y):x \in (x_0 + R/4, x_0 + 3R/4], y \in (y_0 + R/4, y_0 + 3R/4]\}} \{h(x, y) - f(x)g(y)\} dx dy \right]^2 \\ &\geq c^2 \int_{\{(x,y):x \in (x_0 + R/4, x_0 + 3R/4], y \in (y_0 + R/4, y_0 + 3R/4]\}} dx dy = c^2 R^2 / 4, \end{aligned}$$

where the first inequality follows since $h(x, y) - f(x)g(y) > 0$ in \mathcal{A} , and the second inequality follows since the minimum value is $c > 0$. Therefore, it follows that $\frac{1}{N} \frac{(o_C - e_C)^2}{e_C}$ converges uniformly almost surely to a positive constant greater than $c' = c^2 R^2 / 4$,

$$Pr \left(\lim_{N \rightarrow \infty} \frac{1}{N} \frac{(o_C - e_C)^2}{e_C} \geq c' \right) = 1. \quad (27)$$

The partition \mathcal{I} either contains the cell C , or a group of cells that divide C . By Jensen’s inequality, it follows that in the latter case the score is made larger, see the arguments leading to expression

(17). It thus follows that the score $T^{\mathcal{I}}/N$ converges uniformly almost surely to a positive constant greater than c' .

Let $|\Gamma|$ denote the number of $\Gamma\{(x_1, y_1), \dots, (x_N, y_N)\}$. Since $\mathcal{I} \in \Gamma\{(x_1, y_1), \dots, (x_N, y_N)\}$ was arbitrarily fixed, it follows that $\frac{1}{|\Gamma|} \sum_{\mathcal{I} \in \Gamma} \frac{T^{\mathcal{I}}}{N}$ converges almost surely to a positive constant greater than $c'/2$. Since $S_{m \times m} \geq \sum_{\mathcal{I} \in \Gamma} T^{\mathcal{I}}$, it follows that almost surely,

$$\lim_{N \rightarrow \infty} \frac{S_{m \times m}}{|\Gamma|N} \geq \frac{c'}{2}. \quad (28)$$

We shall show that $\lim_{N \rightarrow \infty} |\Gamma|/\{\binom{N-1}{m-1}\binom{N-1}{m-1}\}$ is bounded below by a positive constant. First, we shall consider the case that m is finite. In this case, a subset of $\Gamma\{(x_1, y_1), \dots, (x_N, y_N)\}$ is the set of all partitions with $m-2$ grid points in \mathcal{A}_1 , and one grid point in \mathcal{A}_2 , for both x -values and y -values. Therefore, $|\Gamma| \geq \binom{N_{1x}}{m-2} N_{2x} \binom{N_{1y}}{m-2} N_{2y}$. Simple algebraic manipulations lead to the following expression for $\frac{\binom{N_{1x}}{m-2} N_{2x}}{\binom{N-1}{m-1}}$:

$$(m-1) \frac{N_{2x}}{N-1} \frac{N_{1x}}{N-2} \cdots \frac{N_{1x}-m+3}{N-1-m+2}.$$

Since N_{1x}/N converges almost surely to $\int_{x_0}^{x_0+R/4} f(x)dx$, and similarly N_{2x}/N converges almost surely to $\int_{x_0+3R/4}^{x_0+R} f(x)dx$, then for $m \geq 3$ finite it follows that $\frac{\binom{N_{1x}}{m-2} N_{2x}}{\binom{N-1}{m-1}}$ converges almost surely to a positive constant. Similarly, $\frac{\binom{N_{1y}}{m-2} N_{2y}}{\binom{N-1}{m-1}}$ converges almost surely to a positive constant. Therefore, $|\Gamma|/\{\binom{N-1}{m-1}\binom{N-1}{m-1}\}$ is bounded away from zero.

Second, we shall consider the case that $m \rightarrow \infty$. The complement of Γ , Γ^C , is the set of contingency tables with no grid point in \mathcal{A}_1 or in \mathcal{A}_2 . An upper bound for $|\Gamma^C|$ is:

$$\binom{N-1}{m-1} \left\{ \binom{N-1-N_{1x}}{m-1} + \binom{N-1-N_{2x}}{m-1} + \binom{N-1-N_{1y}}{m-1} + \binom{N-1-N_{2y}}{m-1} \right\}$$

Note that since $|\Gamma| = \{\binom{N-1}{m-1}\binom{N-1}{m-1}\} - |\Gamma^C|$, it is enough to show that $|\Gamma^C|/\{\binom{N-1}{m-1}\binom{N-1}{m-1}\}$ converges to zero as $N \rightarrow \infty$. Simple algebraic manipulations lead to the following expression for $\frac{\binom{N-1-N_{1x}}{m-1}}{\binom{N-1}{m-1}}$:

$$\left(1 - \frac{N_{1x}}{N-1}\right) \cdots \left(1 - \frac{N_{1x}}{N-1-(m-2)}\right) \leq \left(1 - \frac{N_{1x}}{N-1}\right)^m.$$

Since N_{1x}/N converges almost surely to a positive fraction $\int_{x_0}^{x_0+R/4} f(x)dx$, it follows that $\frac{\binom{N-1-N_{1x}}{m-1}}{\binom{N-1}{m-1}}$ converges almost surely to zero. Similarly, since N_{2x}/N , N_{1y}/N and N_{2y}/N converge almost surely to positive fractions, it follows that respectively, $\frac{\binom{N-1-N_{2x}}{m-1}}{\binom{N-1}{m-1}}$, $\frac{\binom{N-1-N_{1y}}{m-1}}{\binom{N-1}{m-1}}$, and $\frac{\binom{N-1-N_{2y}}{m-1}}{\binom{N-1}{m-1}}$ converge almost surely to zero. Thus $|\Gamma^C|/\{\binom{N-1}{m-1}\binom{N-1}{m-1}\}$ converges almost surely to zero.

Since $\lim_{N \rightarrow \infty} |\Gamma| / \left\{ \binom{N-1}{m-1} \binom{N-1}{m-1} \right\}$ is bounded below by a positive constant, it follows from (28) that almost surely,

$$\lim_{N \rightarrow \infty} \frac{S_{m \times m}}{\binom{N-1}{m-1} \binom{N-1}{m-1} N} \geq c'', \quad (29)$$

for some constant $c'' > 0$.

Consider now a random permutation $(\pi y_1), \dots, (\pi y_N)$ of the y -values y_1, \dots, y_N . Let $S_{m \times m}^\pi$ be the test statistic that is computed from the data $(x_1, \pi y_1), \dots, (x_N, \pi y_N)$. By Markov's inequality,

$$\begin{aligned} Pr \left(S_{m \times m}^\pi \geq c'' \binom{N-1}{m-1} \binom{N-1}{m-1} N \mid \vec{x}, \vec{y} \right) &\leq \frac{E(S_{m \times m}^\pi \mid \vec{x}, \vec{y})}{c'' \binom{N-1}{m-1} \binom{N-1}{m-1} N} \\ &\approx \frac{\binom{N-1}{m-1} \binom{N-1}{m-1} (m-1)^2}{c'' \binom{N-1}{m-1} \binom{N-1}{m-1} N}, \end{aligned} \quad (30)$$

where $\vec{x} = (x_1, \dots, x_N)$ and $\vec{y} = (y_1, \dots, y_N)$. The approximation in (30) becomes more accurate the larger N is, since each of the contingency tables is approximately χ^2 with $(m-1)^2$ degrees of freedom. The right hand side of equation (30) goes to zero as $N \rightarrow \infty$, as long as $\lim_{N \rightarrow \infty} \frac{m}{\sqrt{N}} = 0$. Thus,

$$\lim_{N \rightarrow \infty, m/\sqrt{N} \rightarrow 0} Pr \left(S_{m \times m}^\pi \geq c'' \binom{N-1}{m-1} \binom{N-1}{m-1} N \mid \vec{x}, \vec{y} \right) = 0. \quad (31)$$

We now have all the necessary results to complete the proof. Specifically,

$$\lim_{N \rightarrow \infty} Pr(S_{m \times m} \leq S_{1-\alpha}^{tab}) \leq Pr \left(S_{m \times m}^\pi \leq c'' \binom{N-1}{m-1} \binom{N-1}{m-1} N \right) = 0, \quad (32)$$

where the inequality follows from (31), since $S_{1-\alpha}^{tab}$ is below $c'' \binom{N-1}{m-1} \binom{N-1}{m-1} N$ for N large enough, and the equality follows from (29), thus proving item 1 of Theorem 1 for the ADP summation statistic. ■

Appendix B. Proof of Theorem 4

Proof We want to show that for all $\epsilon > 0$, $\lim_{N \rightarrow \infty} Pr \left(\left| \frac{S_{m \times m}^{ADP(L)}}{N|\Pi|} - I_{XY} \right| > \epsilon \right) = 0$ if $\lim_{N \rightarrow \infty} \frac{m}{\sqrt{N}} = 0$ and $\lim_{N \rightarrow \infty} m = \infty$, where $|\Pi| = \left(\binom{N-1}{m-1} \right)^2$ is the number of partitions.

For continuous marginals, the copula function of the joint distribution of (X, Y) is unique, denote it by $c(u, v)$. The mutual information is the negative copula entropy, $H_{UV} = - \int c(u, v) \log c(u, v) dudv$,

$$\begin{aligned} I_{XY} &= \int c(F_X(x), F_Y(y)) f(x) g(y) \log c(F_X(x), F_Y(y)) dx dy \\ &= \int c(u, v) \log c(u, v) dudv = -H_{UV}. \end{aligned} \quad (33)$$

Consider an arbitrary fixed partition $\mathcal{I} = \{(i_1, j_1), \dots, (i_{m-1}, j_{m-1})\} \subset \{1.5, \dots, N - 0.5\}^2$. Recall that $\mathcal{C}(\mathcal{I})$ is the set of $m \times m$ cells that are defined by the partition. For a cell C , let $r_l(C)$ and $r_h(C)$ be, respectively, the lowest and highest x -grid integer values in C . Similarly, let $s_l(C)$ and $s_h(C)$ be, respectively, the lowest and highest y -grid integer values in C .

The entropy of the partition \mathcal{I} is

$$H_{UV}^{\mathcal{I}} = - \sum_{C \in \mathcal{C}(\mathcal{I})} Pr \left(\frac{r_l(C)}{N} \leq U \leq \frac{r_h(C)}{N}, \frac{s_l(C)}{N} \leq V \leq \frac{s_h(C)}{N} \right) \log \left\{ Pr \left(\frac{r_l(C)}{N} \leq U \leq \frac{r_h(C)}{N}, \frac{s_l(C)}{N} \leq V \leq \frac{s_h(C)}{N} \right) \right\}.$$

The corresponding empirical (plug in) estimator is

$$\hat{H}_{UV}^{\mathcal{I}} = - \sum_{C \in \mathcal{C}(\mathcal{I})} \frac{o_C}{N} \log \left(\frac{o_C}{N} \right), \quad o_C = \sum_{i=1}^N I(r_l(C) \leq r_i \leq r_h(C), s_l(C) \leq s_i \leq s_h(C)).$$

Let $H_U^{\mathcal{I}}$ and $H_V^{\mathcal{I}}$ be the fixed marginal entropies of the partition \mathcal{I} :

$$H_U^{\mathcal{I}} = - \sum_{C_x \in \mathcal{C}_x(\mathcal{I})} \frac{r_h(C_x) - r_l(C_x)}{N} \log \left(\frac{r_h(C_x) - r_l(C_x)}{N} \right),$$

$$H_V^{\mathcal{I}} = - \sum_{C_y \in \mathcal{C}_y(\mathcal{I})} \frac{r_h(C_y) - r_l(C_y)}{N} \log \left(\frac{r_h(C_y) - r_l(C_y)}{N} \right),$$

where $\mathcal{C}_x(\mathcal{I})$ and $\mathcal{C}_y(\mathcal{I})$ are the intervals induced by \mathcal{I} in x and in y , respectively. Note that given \mathcal{I} , the observed and expected margins of the partitions are fixed, and therefore

$$H_U^{\mathcal{I}} = - \sum_{C \in \mathcal{C}(\mathcal{I})} o_C \log \left(\frac{r_h(C) - r_l(C)}{N} \right) \quad (34)$$

$$= - \sum_{C \in \mathcal{C}(\mathcal{I})} Pr(r_l(C) \leq U \leq r_h(C), s_l(C) \leq V \leq s_h(C)) \log \left(\frac{r_h(C) - r_l(C)}{N} \right) \quad (35)$$

$$H_V^{\mathcal{I}} = - \sum_{C \in \mathcal{C}(\mathcal{I})} o_C \log \left(\frac{s_h(C) - s_l(C)}{N} \right) \quad (36)$$

$$= - \sum_{C \in \mathcal{C}(\mathcal{I})} Pr(r_l(C) \leq U \leq r_h(C), s_l(C) \leq V \leq s_h(C)) \log \left(\frac{s_h(C) - s_l(C)}{N} \right) \quad (37)$$

The following simple derivation shows that the likelihood ratio score $T^{\mathcal{I}}$ is a linear combination of $\hat{H}_{UV}^{\mathcal{I}}$, $H_U^{\mathcal{I}}$ and $H_V^{\mathcal{I}}$:

$$\begin{aligned} T^{\mathcal{I}} &= \sum_{C \in \mathcal{C}(\mathcal{I})} o_C \log \frac{o_C}{e_C} \\ &= \sum_{C \in \mathcal{C}(\mathcal{I})} o_C \log \frac{o_C}{N} - \sum_{C \in \mathcal{C}(\mathcal{I})} o_C \log \frac{e_C}{N} \\ &= -N \hat{H}_{UV}^{\mathcal{I}} - \sum_{C \in \mathcal{C}(\mathcal{I})} o_C \log \left(\frac{N r_h(C) - r_l(C)}{N} \frac{s_h(C) - s_l(C)}{N} \frac{1}{N} \right) \\ &= -N \hat{H}_{UV}^{\mathcal{I}} + N H_U^{\mathcal{I}} + N H_V^{\mathcal{I}}, \end{aligned}$$

where the last equality follows from equations (34) and (36).

Let $E(\cdot)$ denote the expectation of a random variable. We bound from above our probability of interest by a sum of three probabilities as follows.

$$\begin{aligned}
 Pr\left(\left|\frac{S_{m \times m}^{ADP}}{N|\Pi|} - I_{XY}\right| > \epsilon\right) &= Pr\left(\left|\frac{\sum_{\mathcal{I}} T_{m \times m}^{\mathcal{I}}(L)}{N|\Pi|} - I_{XY}\right| > \epsilon\right) \\
 &= Pr\left(\left|\frac{\sum_{\mathcal{I}} (H_U^{\mathcal{I}} + H_V^{\mathcal{I}} - \hat{H}_{UV}^{\mathcal{I}})}{|\Pi|} + H_{UV}\right| > \epsilon\right) \\
 &= Pr(|\sum_{\mathcal{I}} (H_U^{\mathcal{I}} + H_V^{\mathcal{I}} - \hat{H}_{UV}^{\mathcal{I}} + H_{UV})| > |\Pi|\epsilon) \\
 &= Pr(|\sum_{\mathcal{I}} (-\hat{H}_{UV}^{\mathcal{I}} + E(\hat{H}_{UV}^{\mathcal{I}})) + \sum_{\mathcal{I}} (-E(\hat{H}_{UV}^{\mathcal{I}}) + H_{UV}^{\mathcal{I}}) \\
 &\quad + \sum_{\mathcal{I}} (-H_{UV}^{\mathcal{I}} + H_U^{\mathcal{I}} + H_V^{\mathcal{I}} + H_{UV})| > |\Pi|\epsilon) \\
 &\leq Pr(|\sum_{\mathcal{I}} (-\hat{H}_{UV}^{\mathcal{I}} + E(\hat{H}_{UV}^{\mathcal{I}}))| + |\sum_{\mathcal{I}} (-E(\hat{H}_{UV}^{\mathcal{I}}) + H_{UV}^{\mathcal{I}})| \\
 &\quad + |\sum_{\mathcal{I}} (-H_{UV}^{\mathcal{I}} + H_U^{\mathcal{I}} + H_V^{\mathcal{I}} + H_{UV})| > |\Pi|\epsilon) \\
 &\leq Pr(|\sum_{\mathcal{I}} (-\hat{H}_{UV}^{\mathcal{I}} + E(\hat{H}_{UV}^{\mathcal{I}}))| > |\Pi|\epsilon/3) \tag{38}
 \end{aligned}$$

$$+ Pr(|\sum_{\mathcal{I}} (-E(\hat{H}_{UV}^{\mathcal{I}}) + H_{UV}^{\mathcal{I}})| > |\Pi|\epsilon/3) \tag{39}$$

$$+ Pr(|\sum_{\mathcal{I}} (-H_{UV}^{\mathcal{I}} + H_U^{\mathcal{I}} + H_V^{\mathcal{I}} + H_{UV})| > |\Pi|\epsilon/3), \tag{40}$$

where the last inequality follows from $\{|\sum_{\mathcal{I}} (-\hat{H}_{UV}^{\mathcal{I}} + E(\hat{H}_{UV}^{\mathcal{I}}))| + |\sum_{\mathcal{I}} (-E(\hat{H}_{UV}^{\mathcal{I}}) + H_{UV}^{\mathcal{I}})| + |\sum_{\mathcal{I}} (-H_{UV}^{\mathcal{I}} + H_U^{\mathcal{I}} + H_V^{\mathcal{I}} + H_{UV})| > |\Pi|\epsilon\} \subseteq \{(|\sum_{\mathcal{I}} (-\hat{H}_{UV}^{\mathcal{I}} + E(\hat{H}_{UV}^{\mathcal{I}}))| > |\Pi|\epsilon/3) \cup (|\sum_{\mathcal{I}} (-E(\hat{H}_{UV}^{\mathcal{I}}) + H_{UV}^{\mathcal{I}})| > |\Pi|\epsilon/3) \cup |\sum_{\mathcal{I}} (-H_{UV}^{\mathcal{I}} + H_U^{\mathcal{I}} + H_V^{\mathcal{I}} + H_{UV})| > |\Pi|\epsilon/3\}$ and Bonferroni's inequality.

We will show that the three probabilities (38)–(40) vanish as $N \rightarrow \infty, m \rightarrow \infty, \frac{m}{\sqrt{N}} \rightarrow 0$, thus proving the theorem.

The probability (38) can be upper-bounded as follows,

$$\begin{aligned}
 Pr\left(|\sum_{\mathcal{I}} (\hat{H}_{UV}^{\mathcal{I}} - E(\hat{H}_{UV}^{\mathcal{I}}))| \geq |\Pi|\epsilon/3\right) &\leq \sum_{\mathcal{I}} Pr\left(|\hat{H}_{UV}^{\mathcal{I}} - E(\hat{H}_{UV}^{\mathcal{I}})| \geq \epsilon/3\right) \\
 &\leq |\Pi| 3e^{-\frac{N}{2}\epsilon^2/9\frac{1}{(\log N)^2}} \tag{41}
 \end{aligned}$$

where the first inequality follows from the fact that $\{|\sum_{\mathcal{I}} (\hat{H}_{UV}^{\mathcal{I}} - E(\hat{H}_{UV}^{\mathcal{I}}))| \geq |\Pi|\epsilon/3\} \subseteq \{\cup_{\mathcal{I} \in \Pi} (\hat{H}_{UV}^{\mathcal{I}} - E(\hat{H}_{UV}^{\mathcal{I}})) \geq \epsilon/3\}$ and Bonferroni's inequality, and the second inequality follows from the upper bound (3.4) in Paninski (2003) for the plug in estimator for a given partition \mathcal{I} . This probability goes to zero as $N \rightarrow \infty$, since $|\Pi|$ is $O(N^{\sqrt{N}})$ and $\lim_{N \rightarrow \infty} O(N^{\sqrt{N}})e^{-\frac{N}{2}\epsilon^2/9\frac{1}{(\log N)^2}} = 0$.

The event in the second probability (39) is not random, so we need to show that $|\sum_{\mathcal{I}}(E(\hat{H}_{UV}^{\mathcal{I}}) - H_{UV}^{\mathcal{I}})| < |\Pi|\epsilon/3$ for N large enough. Proposition 1 in Paninski (2003) states that $0 \leq (H_{UV}^{\mathcal{I}} - E(\hat{H}_{UV}^{\mathcal{I}})) \leq \log(1 + \frac{(m-1)^2-1}{N})$. Therefore,

$$|\sum_{\mathcal{I}}(E(\hat{H}_{UV}^{\mathcal{I}}) - H_{UV}^{\mathcal{I}})| \leq |\Pi| \log(1 + \frac{(m-1)^2-1}{N})$$

Clearly, the RHS is below $|\Pi|\epsilon/3$ for N large enough, if $\lim_{N \rightarrow \infty} \frac{m}{\sqrt{N}} = 0$.

It remains to show that (40) vanishes as $N \rightarrow \infty$. This event is not random, so we will show that

$$\lim_{N \rightarrow \infty} \frac{|\sum_{\mathcal{I}}(-H_{UV}^{\mathcal{I}} + H_U^{\mathcal{I}} + H_V^{\mathcal{I}} + H_{UV})|}{|\Pi|} < \epsilon/3.$$

By the mean value theorem, for cell C there exists a point (u_C, v_C) in C such that

$$Pr(r_l(C) \leq U \leq r_h(C), s_l(C) \leq V \leq s_h(C)) = c(u_C, v_C) \frac{r_h(C) - r_l(C)}{N} \frac{s_h(C) - s_l(C)}{N}.$$

Therefore,

$$\begin{aligned} & -H_{UV}^{\mathcal{I}} \\ &= \sum_{C \in \mathcal{C}(\mathcal{I})} c(u_C, v_C) \frac{r_h(C) - r_l(C)}{N} \frac{s_h(C) - s_l(C)}{N} \log \left(c(u_C, v_C) \frac{r_h(C) - r_l(C)}{N} \frac{s_h(C) - s_l(C)}{N} \right) \\ &= \sum_{C \in \mathcal{C}(\mathcal{I})} \frac{r_h(C) - r_l(C)}{N} \frac{s_h(C) - s_l(C)}{N} c(u_C, v_C) \log c(u_C, v_C) \\ &+ \sum_{C \in \mathcal{C}(\mathcal{I})} Pr(r_l(C) \leq U \leq r_h(C), s_l(C) \leq V \leq s_h(C)) \log \frac{r_h(C) - r_l(C)}{N} \\ &+ \sum_{C \in \mathcal{C}(\mathcal{I})} Pr(r_l(C) \leq U \leq r_h(C), s_l(C) \leq V \leq s_h(C)) \log \frac{s_h(C) - s_l(C)}{N} \\ &= \sum_{C \in \mathcal{C}(\mathcal{I})} \frac{r_h(C) - r_l(C)}{N} \frac{s_h(C) - s_l(C)}{N} c(u_C, v_C) \log c(u_C, v_C) - H_U^{\mathcal{I}} - H_V^{\mathcal{I}}, \end{aligned} \tag{42}$$

where the last equality follows from equations (35) and (37).

By the definition of the Riemann integral, expression (42) can be made arbitrarily close to $-H_{UV}$. Specifically, there exists a $0 < d(\epsilon) < 1$ such that if all cells satisfy $\frac{r_h(C) - r_l(C)}{N} < d$ and $\frac{s_h(C) - s_l(C)}{N} < d$, then

$$\left| \sum_{C \in \mathcal{C}(\mathcal{I})} \frac{r_h(C) - r_l(C)}{N} \frac{s_h(C) - s_l(C)}{N} c(u_C, v_C) \log c(u_C, v_C) + H_{UV} \right| < \epsilon/3.$$

Therefore, it follows that for any partition $\mathcal{I} \in \Pi$ for which all cells satisfy $\frac{r_h(C) - r_l(C)}{N} < d$ and $\frac{s_h(C) - s_l(C)}{N} < d$, then we have

$$|(-H_{UV}^{\mathcal{I}} + H_U^{\mathcal{I}} + H_V^{\mathcal{I}} + H_{UV})| < \epsilon/3.$$

It remains to show that the contribution of the fraction of partitions that do not satisfy $\frac{r_h(C) - r_l(C)}{N} < d$ and $\frac{s_h(C) - s_l(C)}{N} < d$ goes to zero as $N \rightarrow \infty$. Since the probability of selecting an x -value (or y -value) for a partition that will have a cell larger than d is $1 - d$, the fraction of “bad” partitions is upper-bounded by

$$\frac{2m \binom{N(1-d)}{m-1} \binom{N}{m-1}}{\left(\binom{N-1}{m-1}\right)^2} \leq 2m(1-d)^{m-1}.$$

Since $m \rightarrow \infty$ the fraction of bad partitions goes to zero.

Note that $|(-H_{UV}^{\mathcal{I}} + H_U^{\mathcal{I}} + H_V^{\mathcal{I}} + H_{UV})|$ is at most $O(\log m^2)$ because by Jensen’s inequality, $|H_{UV}^{\mathcal{I}}| \leq \log m^2$, $|H_U^{\mathcal{I}}| \leq \log m^2$, $|H_V^{\mathcal{I}}| \leq \log m^2$, and $|H_{UV}| = I_{XY}$ is assumed to be bounded. Therefore,

$$\frac{|\sum_{\mathcal{I}}(-H_{UV}^{\mathcal{I}} + H_U^{\mathcal{I}} + H_V^{\mathcal{I}} + H_{UV})|}{|\Pi|} \leq \epsilon/3 + O(\log m^2 m(1-d)^{m-1}).$$

Since the second term of the RHS goes to zero as $N \rightarrow \infty$, $m \rightarrow \infty$, $\frac{m}{\sqrt{N}} \rightarrow 0$, the proof is complete. \blacksquare

Appendix C. Proof of Theorem 6

Proof We shall prove items 1 and 2 for the DDP statistic only, since the proof for the ADP statistic is very similar. We shall use the notation of Section A.1. Let \hat{m} be the value of m with minimum p -value,

$$\hat{m} = \arg \min_{m \in \{2, \dots, m_{\max}\}} p_2, \dots, p_m.$$

To prove item 1, we note that from equation (19) it follows that under the alternative,

$$\lim_{N \rightarrow \infty} \frac{S_{\hat{m} \times \hat{m}}}{\binom{N}{\hat{m}-1} (N - \hat{m} + 1)} \geq c''.$$

Under the null,

$$\begin{aligned} & Pr \left(S_{\hat{m} \times \hat{m}}^{\pi} \geq c'' \binom{N}{\hat{m}-1} (N - \hat{m} + 1) \mid \vec{x}, \vec{y} \right) \\ & \leq \sum_{m=2}^{m_{\max}} Pr \left(S_{m \times m}^{\pi} \geq c'' \binom{N}{m-1} (N - m + 1) \mid \vec{x}, \vec{y} \right) \\ & \leq m_{\max} \max_{m \in \{2, \dots, m_{\max}\}} \frac{\binom{N}{m-1} (m-1)^2}{c'' \binom{N}{m-1} (N - m + 1)} \leq m_{\max} \frac{(m_{\max} - 1)^2}{c'' (N - m_{\max} + 1)}, \end{aligned}$$

where the first inequality is the Bonferroni inequality, and the second inequality follows from equation (20). Since the last term goes to 0 as $N \rightarrow \infty$ if $\lim_{N \rightarrow \infty} m/N^{1/3} = 0$, the proof of item 1 is complete.

The proof of item 2 is very similar to the proof in Section A.1, the only modification is an additional application of Bonferroni's inequality under the null:

$$\begin{aligned} Pr_{H_0} \left(M_{\hat{m} \times \hat{m}}^{DDP} \geq \frac{c'}{2}(N - \hat{m} + 1) \right) &\leq \sum_{m=2}^{m_{\max}} Pr_{H_0} \left(M_{m \times m}^{DDP} \geq \frac{c'}{2}(N - m + 1) \right) \\ &\leq m_{\max} \max_{m \in \{2, \dots, m_{\max}\}} \binom{N}{m} e^{-\frac{c'}{8}(N-m+1)-(m-1)^2/4} \\ &\leq m_{\max} \binom{N}{m_{\max}} e^{-\frac{c'}{8}(N-m_{\max}+1)-(m_{\max}-1)^2/4}, \end{aligned}$$

where the first inequality in the last row follows from (24). Since $m_{\max} \binom{N}{m_{\max}}$ is at most $O(N \times N^{\sqrt{N}})$, and since

$$e^{-\frac{c'}{8}(N-m_{\max}+1)-(m_{\max}-1)^2/4} = O(e^{-\frac{c'}{8}N}),$$

it follows that the last expression goes to zero as $N \rightarrow \infty$. ■

Appendix D. An Example of Mutual Information Estimation

We examined our suggested estimator in the following setup. For $m = 15$ and $N \in \{300, 1000\}$ sample points drawn from a two-component Gaussian mixture, we simulated 50 data sets and computed the estimated mutual information using $S_{m \times m}^{DDP}$, $S_{m \times m}^{ADP}$, and the histogram estimator. The Gaussian mixture density was

$$0.8 \times f_{\mathcal{N}} \left(\begin{pmatrix} 0.5 \\ 0.5 \end{pmatrix}, \begin{pmatrix} 0.05 & 0.025 \\ 0.025 & 0.05 \end{pmatrix} \right) + 0.2 \times f_{\mathcal{N}} \left(\begin{pmatrix} 0.125 \\ 0.675 \end{pmatrix}, \begin{pmatrix} 0.01 & 0 \\ 0 & 0.01 \end{pmatrix} \right),$$

where $f_{\mathcal{N}}(\mu, \Sigma)$ is the bivariate normal distribution with mean μ and covariance matrix Σ . For each partition, we applied the Miller–Madow correction (Paninski, 2003), a simple and well-known modification that estimates the systematic error of the histogram estimators and improves the finite-sample properties of these estimators. We compared it to the histogram estimator, as well as to the estimator based on the DDP statistic that considers only a subset of all possible partitions. Table 5 shows that the variability and the bias decrease as N increases for all estimators, and that the ADP estimator is the least variable, as is intuitively expected since it is the average over many partitions. In practice, it is difficult to identify the optimal m : it should not be too small so that the local dependence structure is not missed, causing large bias, nor should it be too large so that the grid created is too sparse, causing large variance.

Appendix E. Algorithm for the DDP Statistic

For the DDP statistic, the algorithm is very similar to that for the ADP statistic, except that only DDP are considered on the grid of ranked data $\{1, \dots, N\}^2$. The algorithm is slightly more complex because the number of partitions that include a cell C depends on the partition size m , on the data, and on the type of cell, with four possible types. Specifics follow for internal cells. The first type is a cell C for which there is a sample point that falls on the boundary of C but not on one of its

	N = 300	N = 1000
Histogram	0.3165 (0.0052)	0.1854 (0.0029)
Data derived partitions	0.3010 (0.0030)	0.1879 (0.0022)
All derived partitions	0.2954 (0.0028)	0.1860 (0.0021)

Table 5: The average (standard error) of the mutual information estimates using random samples of size 300 (column 2) and 1000 (column 3) from the two component Gaussian mixture, by the following methods: the naive histogram estimator that partitions each axis to 15 intervals of equal count, $S_{15 \times 15}^{ADP}/(N|\Pi_{15}^{ADP}|)$, and $S_{15 \times 15}^{DDP}/(N|\Pi_{15}^{DDP}|)$. The true mutual information value was 0.1784.

corners. Then no DDP can ever have C as a cell, and therefore the number of DDP that include C is zero. For example, in Figure 7 (middle panel), if the open circle is an observation, and therefore the filled circle with the same y value is not, since there are no ties, then any DDP with this y value will necessarily partition C at the x value of the open circle observation, and thus C cannot be a cell in any DDP. For the remaining three types of cells, if there are sample points that fall on the boundary of C they are necessarily on the corners of C . These types of cells differ by the number of observations that determine the cell. The second type is a cell C defined by two observed points. Then, the number of DDP that include C is the number of ways to choose $m - 3$ points from the points in the four outer areas defined by $(0, r_l) \times (0, s_l)$, $(0, r_l) \times (s_h, N]$, $(r_h, N] \times (0, s_l)$, and $(r_h, N] \times (s_h, N]$, see Figure 7 (left panel) for illustration. The number of points in the four areas is calculable in $O(1)$ using A , as defined in equation (10). Specifically, the count of samples that fall strictly inside any cell with rank ranges $r \in [r_l, r_h]$ and $s \in [s_l, s_h]$ is:

$$o_C = A(r_h - 1, s_h - 1) - A(r_l, s_h - 1) - A(r_h - 1, s_l) + A(r_l, s_l).$$

The third and fourth type are cells defined by three or four observed points, respectively, see Figure 7 (middle and right panels). Now the number of DDP that include C is exactly the number of ways to choose $m - 4$ and $m - 5$ points, respectively, from the points in the four outer areas. Again, the number of points in the four outer areas is calculable in $O(1)$ using A . Since all cells are defined by two, three, or four points, there are no additional types of cells.

Let us denote the number of sample points in the four outer areas of a cell C by OUT . Further denote by $C(OUT)$ the group of all cells which have exactly OUT points in the outer areas. For the sake of brevity, let's consider only cells of type 2. Similarly to equation (12) for the ADP statistic, the contribution to the score $S_{m \times m}^{DDP}$ of internal cells of type 2 can be written as

$$\sum_{OUT=2}^{N-2} n(OUT, m) \sum_{C \in C(OUT)} t_C = \sum_{OUT=2}^{N-2} n(OUT, m) T(OUT). \quad (44)$$

The algorithm proceeds as follows. First in a preprocessing phase we perform two computations: 1) go over all cells and calculate t_C and OUT for each cell and update $T(OUT) = T(OUT) + t_C$. This stage takes $O(N^4)$; 2) for all $u, v \in \{0, \dots, N\}$ we calculate and store all $\binom{u}{v}$ in $O(N^2)$ steps using Pascal's triangle method.




Figure 7: Inner cell example C (orange), and the sample points that define a partition where C is a cell (green). If any of the sample points ranked on a boundary column or row of C is not on the corner of C , then C can never be a DDP cell (offending point in purple, middle panel). An inner cell can be defined either by two data points (left), by three data points (middle), or by four data points (right).

Now for each m , since $n(OUT, m) = \binom{OUT}{m-3}$, given $T(OUT)$, clearly equation (44) can be calculated in $O(N)$ for a fixed m and in $O(N^2)$ for all m s. Therefore the total complexity is again $O(N^4)$ due to the preprocessing phase.

Appendix F. A Fast Algorithm for Computing the HHG Statistic

Here we describe a fast algorithm for computing the univariate original (distribution-dependent) HHG test statistic.

If $D_X(x_1, x_2)$ and $D_Y(y_1, y_2)$ are distance metrics in the variables tested for independence (for which, again, we have a paired sample with N i.i.d samples), the HHG test requires computing $N(N-1)$ different 2×2 contingency tables according to the following partitions of the distance-distance plane. For every “origin” sample i , and for every “radius” sample j , the remaining $N-2$ samples are classified according to whether their X and Y distances from i are both smaller than the X and Y distances from j to i , or if only X , only Y , or neither are smaller than the respective j to i distances. In the general case, generating contingency tables for all pairs can be done in $N^2 \log N$, as described in (Heller et al., 2013). In the univariate case, an $O(N^2)$ algorithm proceeds as follows.

Instead of working in the distance-distance plane, the algorithm is specified in terms of the (x, y) , i.e., the sample plane. It is sufficient to consider the discrete grid expanded from unique X samples and unique Y samples actually observed (these can be identified in $O(N \log N)$). The double cumulative sum over this $N \times N$ grid is computed as in Section 3.1 of the main text, with the only difference being that after A is initialized to all zeros, it is updated sequentially with $A(r_i, s_i) = A(r_i, s_i) + 1$ for every sample of ranks in x and ranks in y , (r_i, s_i) , $i = 1, 2, \dots, N$, to account for possible ties.

Since, in the univariate case, $D_X(x_i, x_k) < D_X(x_i, x_j)$ is equivalent to $|x_k - x_i| < |x_j - x_i|$, and similarly for y , partition cells are simply axis-aligned rectangles, as in the distribution-free test. Here, however, only sample j is a vertex, and sample i is the center of mass. The diagonal-opposing vertex from j may not even be a point in the sample, and thus is not directly captured by A . Still, the appropriate point to sample A in, for computing the contingency table cell in $O(1)$, can be found in $O(1)$ additional amortized time, as follows:

1. Sort once the unique values of x , and do the same for y .
2. When traversing all pairs, first traverse i .
 - (a) For every i , traverse the sorted values of x with two concurrent iterators starting from x_i , one moving right (i.e., from low to high x) and one advancing left.
 - (b) Subsequently to each step taken with the right iterator, arriving at an x_j , advance the left iterator until a value is encountered which is farther from x_i than x_j is, and this is the opposing vertex x coordinate value for the rectangle for the pair i, j .

The process, which is depicted in Figure 8, takes $O(N)$ time for N values, and is repeated for the y axis.




Figure 8: Finding the grid coordinates to sample for the $O(N^2)$ univariate (distribution-dependent) HHG algorithm.

Appendix G. Additional Simulations for the Two-Sample Problem

Table 6 shows a comparison of the performance of the minimum p -value and Fisher-combined test statistics. The variant with highest power in most setups (specifically, setups 2,4,5,6,8,9,10,11,12) is the Fisher-combined p -value using aggregation by summation and $m_{\max} = 50$. However, the Fisher-combined p -value test is sensitive to the choice of m_{\max} , which is unknown in practice. Since we view this as a significant weakness of the Fisher-combined statistic, and since the minimum p -value does not have this weakness and has very good power when compared with Fisher as well as when compared with other tests (with a large range of m_{\max} values examined), we recommend the minimum p -value test statistic.

Table 7 shows the power using different priors for regularization. The priors are as follows: $\pi(m) = \sqrt{N}^m \exp(-\sqrt{N})/m!$ for Poisson; $\pi(m) = \binom{N-1}{m-1} p^m (1-p)^{(N-m)}$ for Binomial, with $p = 0.119$ so that the penalty becomes that of the Akaike information criterion (i.e., $\log \pi(\mathcal{I}|m)\pi(m) = -2m$); $\pi(m) = 1/K$ for fixed K for Uniform. We also considered the prior in Jiang et al. (2014), resulting in the DS test with penalty $-\lambda_0 \log N(m-1)$. According to the recommendation in Jiang et al. (2014), we chose the value of λ_0 so that the level under the null is as close as possible to 0.05 from below, so $\lambda_0 = 1.11088$ for $N = 100$, and $\lambda_0 = 0.904$ for $N = 500$. The Poisson prior (with rate \sqrt{N}) was by far the best among all priors considered, and its power was comparable to that of the minimum p -value displayed in Table 6.

For sample size $N = 100$, we examined the distributions depicted in Figure 10, and we used 20,000 simulated data sets, in each of the configurations. Table 9 and Figure 11 show the power for

Setup	Minimum p -value				Fisher combined p -value			
	Max aggregation, m_{\max}		Sum aggregation, m_{\max}		Max aggregation, m_{\max}		Sum aggregation, m_{\max}	
	= 50	= 149	= 50	= 149	= 50	= 149	= 50	= 149
1	0.836	0.825	0.500	0.491	0.631	0.424	0.549	0.560
2	0.810	0.799	0.883	0.873	0.819	0.701	0.917	0.898
3	0.783	0.785	0.553	0.733	0.867	0.847	0.502	0.782
4	0.836	0.827	0.945	0.937	0.749	0.567	0.952	0.882
5	0.607	0.592	0.708	0.686	0.531	0.381	0.760	0.655
6	0.829	0.818	0.836	0.820	0.749	0.567	0.882	0.821
7	0.358	0.339	0.516	0.492	0.201	0.135	0.430	0.287
8	0.770	0.752	0.795	0.775	0.542	0.361	0.814	0.672
9	0.531	0.512	0.641	0.613	0.344	0.227	0.635	0.470
10	0.728	0.711	0.824	0.806	0.569	0.393	0.849	0.721
11	0.556	0.540	0.707	0.686	0.509	0.380	0.756	0.657
12	0.402	0.390	0.599	0.577	0.399	0.293	0.655	0.557
13	0.852	0.844	0.777	0.764	0.809	0.663	0.823	0.822

Table 6: Power of the minimum p -value test statistic (columns 2–5) as well as the Fisher-combined p -value test statistic (columns 6–9), using the different aggregation methods and two values of m_{\max} , for $N = 500$ and the setups of Figure 2. The difference between $m_{\max} = 50$ and $m_{\max} = 149$ is much larger for the Fisher-combined than for the minimum p -value test statistic when aggregated by summation (columns 8–9 versus columns 4–5). Moreover, the power is typically lower for the Fisher-combined than for the minimum p -value test statistic when the aggregation is by maximization (columns 6–7 versus columns 2–3).

Setup	Poisson prior	Poisson prior	Uniform Prior	Binomial Prior	DS prior
	Max aggregation	Sum aggregation	Max aggregation	Max aggregation	Max aggregation
1	0.829	0.597	0.848	0.450	0.845
2	0.855	0.924	0.641	0.745	0.565
3	0.715	0.639	0.282	0.874	0.239
4	0.861	0.920	0.813	0.593	0.794
5	0.659	0.746	0.592	0.411	0.573
6	0.854	0.874	0.804	0.597	0.787
7	0.291	0.310	0.430	0.139	0.432
8	0.719	0.760	0.835	0.371	0.836
9	0.489	0.545	0.599	0.232	0.600
10	0.709	0.793	0.749	0.406	0.744
11	0.607	0.720	0.466	0.390	0.433
12	0.460	0.629	0.298	0.318	0.268
13	0.860	0.857	0.802	0.697	0.776

Table 7: Power using different priors, for $N = 500$ and the setups of Figure 2.

the setups in Figure 10. These results concur with the results in Section 4.1 for $N = 500$, and show that if the number of intersections of the two distributions is at least four, tests statistics with $m \geq 4$ have an advantage.

Appendix H. Simulations of Monotone Relationships

Four monotone relationships are presented in Figure 12. Figure 13 shows that the differences for the summation variants between ADP and DDP are negligible, and that the power decreases with m for all of these variants. Similar conclusions hold for the maximum variants in Table 10. Table 10 further shows that for a linear relationship, Pearson, Spearman, Hoeffding and dCov are clearly

		Minimum p -value							
N	Setup	Max aggregation	Sum Aggregation	Wilcoxon	KS	CVM	AD	HHG	DS
500	Normal shift	0.59	0.82	<u>0.91</u>	0.81	0.89	0.90	0.85	0.73
500	Normal scale	0.76	0.85	0.05	0.36	0.46	0.77	<u>0.88</u>	0.83
500	Normal shift & scale	0.83	0.90	0.46	0.69	0.74	0.88	<u>0.92</u>	0.90
100	Normal shift	0.39	0.58	<u>0.68</u>	0.53	0.65	0.67	0.60	0.49
100	Normal scale	0.51	0.59	0.06	0.20	0.23	0.41	<u>0.63</u>	0.54
100	Normal shift & scale	0.54	0.64	0.49	0.51	0.58	0.64	<u>0.68</u>	0.62

Table 8: Power of competitors (columns 4–9), along with the minimum p -value using the M_m p -values (column 1) and the S_m p -values (column 2).

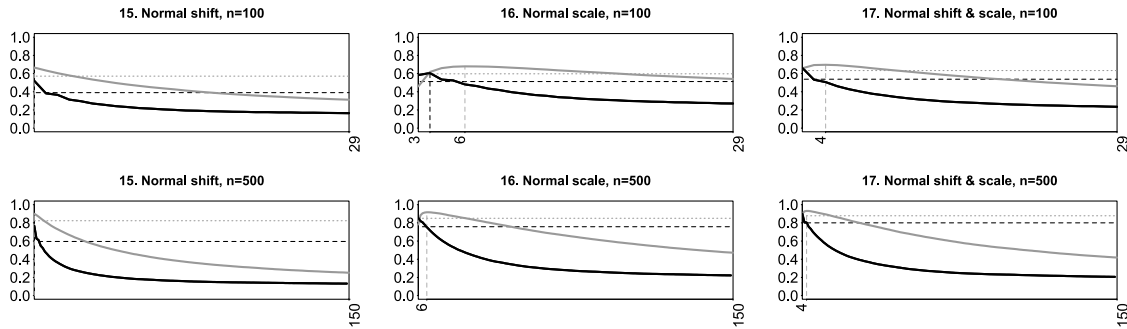


Figure 9: Estimated power for the M_m (black) and S_m (grey) statistics for $m \in \{2, \dots, 149\}$ for the Gaussian shift difference (first column, $N(0, 1)$ versus $N(0.5, 1)$ for $N = 100$, and versus $N(0.3, 1)$ for $N = 500$), Gaussian scale difference (second column, $N(0, 1)$ versus $N(0, 0.6^2)$ for $N = 100$, and versus $N(0, 0.75^2)$ for $N = 500$), and for the Gaussian shift and scale difference (third column, $N(0, 1)$ versus $N(0.36, 0.7^2)$ for $N = 100$, and versus $N(0.2, 0.8^2)$ for $N = 500$). The power of the minimum p -value is the horizontal dashed black line when it combines the p -values based on M_m , and the horizontal dotted grey line when it combines the p -values based on S_m . The vertical lines show the optimal m for M_m (grey) and S_m (black).

superior, but for the remaining three monotone relationships $\min_{m \in \{2, \dots, m_{\max}\}} p_m$ has quite good power properties.

Appendix I. Comparisons of DCov/HHG on Ranks and on Data

In Section 4 we only considered distribution-free competitors. Therefore, we applied dCov and HHG on ranks rather than on data, even though they were designed as permutation tests on data. Usually, the computational advantage of performing the tests on $(rank(X), rank(Y))$ instead of on (X, Y) , due to the distribution-free property of the tests on $(rank(X), rank(Y))$, comes at a cost of lower power, as noted by Székely et al. (2007). Table 11 shows a power comparison of these two permutation tests on ranks and on data. The power in most settings is indeed greater when the tests are used on data, and the maximal difference is almost 30% (in the Spiral setting for HHG) in favour of using the data. However, in some settings the power is actually larger for the test on ranks, e.g.,

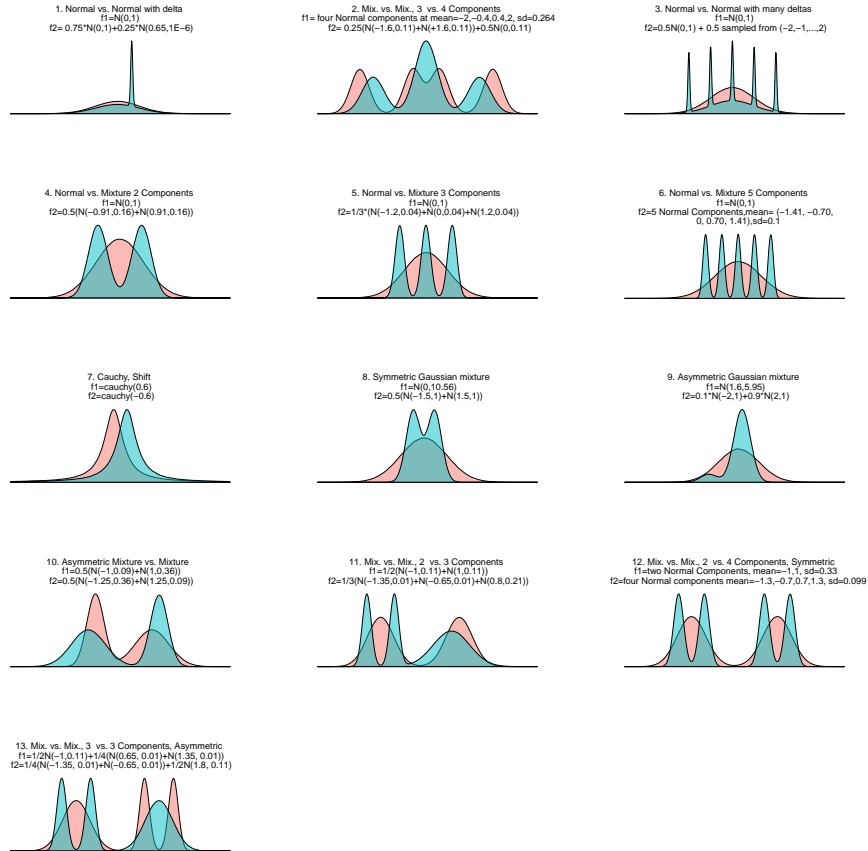


Figure 10: The two distributions in 13 different setups considered for $N = 100$, which differ in the number of intersections of the densities, the range of support where the differences lie, and the whether they are symmetric or not.

DISTRIBUTION-FREE K -SAMPLE AND INDEPENDENCE TESTS

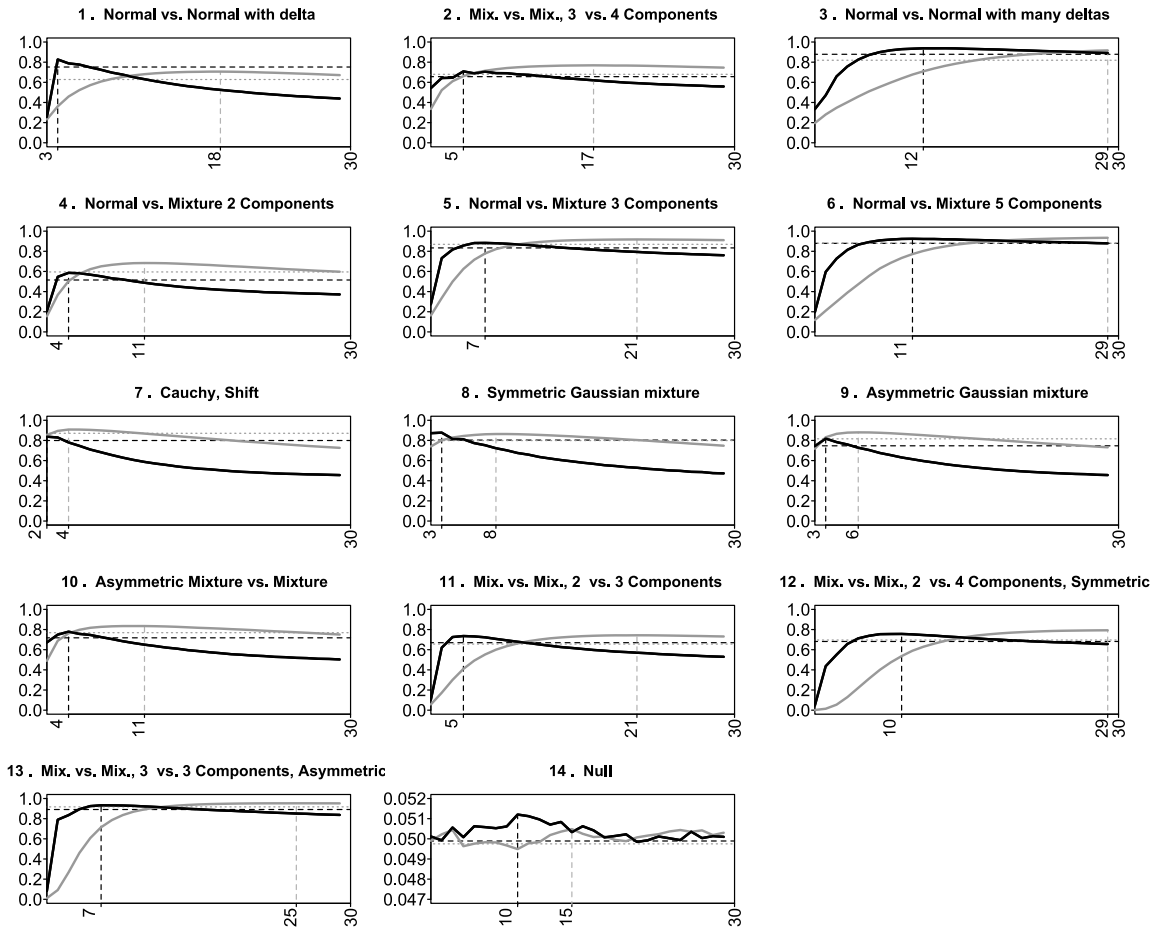


Figure 11: Estimated power with $N = 100$ sample points for the M_m (black) and S_m (grey) statistics for $m \in \{2, \dots, 29\}$ for the setups of Figure 10. The power of the minimum p -value is the horizontal dashed black line when it combines the p -values based on M_m , and the horizontal dotted grey line when it combines the p -values based on S_m . The vertical lines show the optimal m for M_m (black) and S_m (grey).

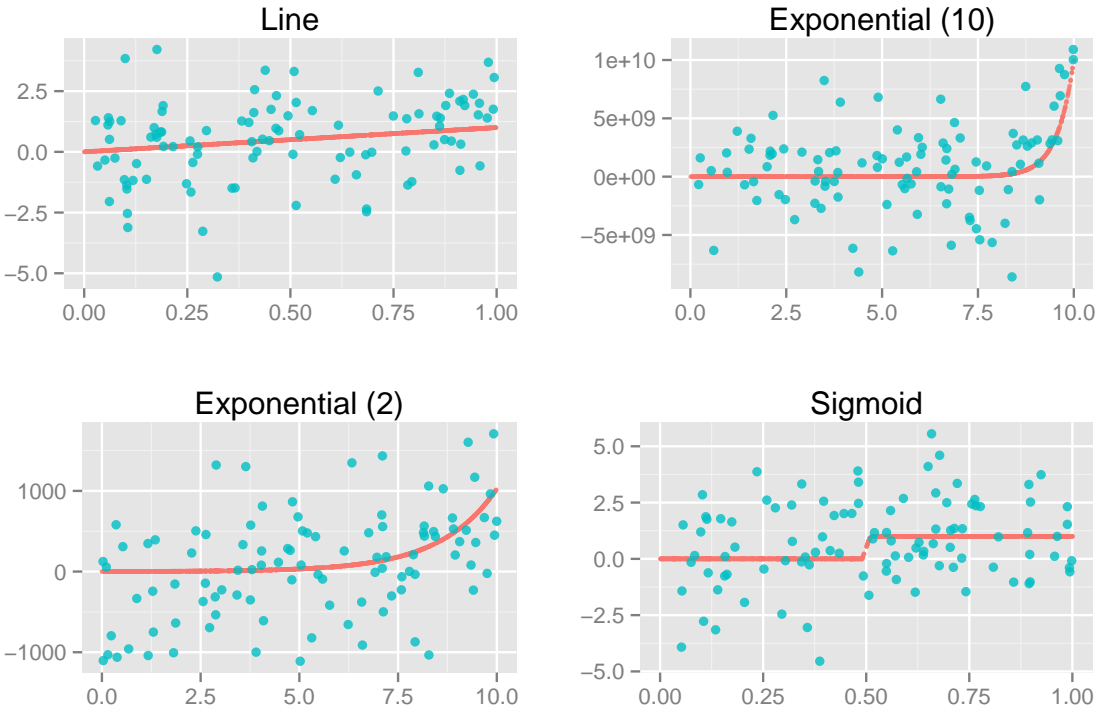


Figure 12: Bivariate monotone relationships (in red), along with a sample of $N = 100$ noisy observations (in blue).




Figure 13: Estimated power for the DDP (black) and ADP (gray) summation variants, using the likelihood ratio score, for the setups from Figure 12, for a sample of size $N = 100$. The horizontal lines are the power of $\min_{m \in \{2, \dots, m_{\max}\}} p_m$ using DDP (dotted black) and using ADP (gray) .

Setup	Min. p -value aggregation								
	by Max	by Sum	Wilcox	KS	CVM	AD	HHG	DS	
1	Normal vs. Normal with delta	<u>0.752</u>	0.628	0.187	0.331	0.294	0.265	0.433	0.660
2	Mix. Vs. Mix., 3 Vs. 4 Components	0.656	<u>0.680</u>	0.000	0.003	0.000	0.063	0.465	0.505
3	Normal vs. Normal with many deltas	<u>0.878</u>	0.819	0.053	0.125	0.100	0.165	0.281	0.342
4	Normal vs. Mixture 2 Components	0.515	<u>0.595</u>	0.052	0.231	0.153	0.154	0.397	0.382
5	Normal vs. Mixture 3 Components	0.834	<u>0.869</u>	0.053	0.177	0.106	0.134	0.317	0.534
6	Normal vs. Mixture 5 Components	0.879	<u>0.880</u>	0.051	0.123	0.084	0.100	0.200	0.373
7	Cauchy, Shift	0.799	0.871	0.847	0.917	0.920	0.893	<u>0.933</u>	0.846
8	Symmetric Gaussian mixture	0.803	0.798	0.035	0.182	0.191	0.491	<u>0.727</u>	<u>0.812</u>
9	Asymmetric Gaussian mixture	0.747	0.816	0.046	0.369	0.407	0.593	<u>0.845</u>	0.740
10	Asymmetric Mixture vs. Mixture	0.718	<u>0.769</u>	0.000	0.157	0.128	0.306	<u>0.655</u>	0.652
11	Mix. Vs. Mix., 2 Vs. 3 Components	<u>0.670</u>	0.656	0.000	0.032	0.011	0.031	0.129	0.441
12	Mix. Vs. Mix., 2 Vs. 4 Comp., Sym.	0.682	<u>0.696</u>	0.000	0.000	0.000	0.000	0.005	0.238
13	Mix. Vs. Mix., 3 Vs. 3 Comp., Asym.	0.891	<u>0.917</u>	0.000	0.000	0.000	0.000	0.053	0.575
14	Null	0.050	0.050	0.049	0.039	0.049	0.050	0.050	0.041

Table 9: Power of competitors (columns 4–9), along with the minimum p -value statistic using the M_m p -values (column 1) and the S_m p -values (column 2), for $N = 100$. The standard error was at most 0.0035. The advantage of the test based on the minimum p -value is large when the number of intersections of the two densities is at least four (setups of rows 2,3, 4,5,6,10,11,12, and 13). The best competitors are HHG and DS, but HHG is essentially an $m \leq 3$ test, and DS penalizes large m s severely, therefore in setups where $m \geq 4$ partitions are better they can perform poorly. Among the two variants in columns 1 and 2, the better choice clearly depends on the range of support in which the differences in distributions occur: aggregation by maximum has better power when the difference between the distributions is very local (setups of rows 1 and 3), and aggregation by summation has better power otherwise. The highest power per row is underlined.

Test	Line	Exp2x	Exp10x	Sigmoid
$\min_{m \in \{2, \dots, 10\}} p_m$ using DDP	0.358	0.763	0.580	0.543
$\min_{m \in \{2, \dots, 10\}} p_m$ using ADP	0.365	0.760	0.555	0.550
Spearman	0.459	0.758	0.396	0.630
Hoeffding	0.446	0.750	0.409	0.637
MIC	0.282	0.198	0.312	0.130
dCov	0.433	0.746	0.395	0.637
HHG	0.337	0.706	0.509	0.545
$M_{2 \times 2}^{DDP}$	0.287	0.688	0.678	0.438
$M_{3 \times 3}^{DDP}$	0.203	0.579	0.569	0.355
$M_{4 \times 4}^{DDP}$	0.177	0.511	0.479	0.301
$M_{2 \times 2}^{ADP}$	0.294	0.715	0.746	0.440

Table 10: The power of competitors (rows 3–7), along with the minimum p -value statistic based on DDP (row 1) and on ADP (row 2) and the different maximum variants (rows 8–11), for $N = 100$. The standard error is at most 0.011.

in the Heavisine and 5Clouds settings for HHG, where the difference is 10% and 16%, respectively, in favour of using the ranked observations. Comparing the power of the HHG and dCov tests on data (in Table 11) to the power of our suggested minimum p -value statistic (in Table 2), we see that although dCov on data may have more power than dCov on ranks, it has far less power than HHG on data, and that HHG on data has less power than our test when the relationship is more complex, especially in the Sine, Heavisine, Spiral and Circles examples.

DISTRIBUTION-FREE K -SAMPLE AND INDEPENDENCE TESTS

Setup	dCov		HHG		$\min_{\{m \in \{2, \dots, 10\}\}} p_m$
	on data	on ranks	on data	on ranks	
W	0.467	0.351	0.876	0.813	0.655
Diamond	0.136	0.076	0.997	0.968	0.919
Parabola	0.418	0.369	0.727	0.795	0.847
2Parabolas	0.183	0.120	0.873	0.726	0.844
Circle	0.001	0.003	0.791	0.858	0.886
Cubic	0.631	0.612	0.660	0.742	0.731
Sine	0.412	0.427	0.803	0.788	0.995
Wedge	0.471	0.327	0.755	0.661	0.595
Cross	0.184	0.138	0.795	0.704	0.704
Spiral	0.182	0.130	0.598	0.335	0.949
Circles	0.054	0.059	0.479	0.357	0.999
Heavisine	0.476	0.470	0.471	0.570	0.710
Doppler	0.747	0.736	0.901	0.914	0.949
5Clouds	0.000	0.001	0.738	0.903	0.996
4Clouds	0.050	0.050	0.050	0.050	0.051

Table 11: Power of the competitors dCov and HHG on ranks as well as on data for $N = 100$ (columns 2-4). The results on ranks are not identical numerically (though very close) to those of Table 2 due to the use of a different seed to generate the data for these same settings. The tests on data have greater power than on ranks in most, but not all, examples. HHG on data has the best power out of these four competitors, but like the HHG on ranks it has a disadvantage in comparison with our novel minimum p -value test when the relationship is more complex, especially in the Sine, Heavisine, Spiral, Circles, and 5Clouds examples (the power for this test was reproduced in column 6 here from Table 2 for convenience).

References

- T.W. Anderson and D.A. Darling. Asymptotic theory of certain "goodness of fit" criteria based on stochastic processes. *The Annals of Mathematical Statistics*, 23(2):193–212, 1952.
- L. Baringhaus and C. Franz. On a new multivariate two-sample test. *Journal of Multivariate Analysis*, 88:190–206, 2004.
- Y. Benjamini and Y. Hochberg. Controlling the false discovery rate - a practical and powerful approach to multiple testing. *J. Roy. Stat. Soc. B Met.*, 57 (1):289–300, 1995.
- J. Blum, J. Kiefer, and M. Rosenblatt. Distribution free tests of independence based on the sample distribution function. *The Annals of Mathematical Statistics*, pages 485–498, 1961.
- K.P. Chwialkowski, D. Sejdinovic, and A. Gretton. A wild bootstrap for degenerate kernel tests. In Z. Ghahramani, M. Welling, C. Cortes, N.d. Lawrence, and K.q. Weinberger, editors, *Advances in Neural Information Processing Systems 27*, pages 3608–3616. Curran Associates, Inc., 2014. URL <http://papers.nips.cc/paper/5452-a-wild-bootstrap-for-degenerate-kernel-tests.pdf>.
- D.A. Darling. The kolmogorov-Smirnov, Cramer-von Mises tests. *The Annals of Mathematical Statistics*, 28(4):823–838, 1957.
- D. Donoho and I. Johnstone. Adapting to unknown smoothness via wavelet shrinkage. *Journal of the American Statistical Association*, 90(432):1200–1224, 1995.
- A. Feuerverger. A consistent test for bivariate dependence. *International Statistical Review*, 61(3): 419–433, 1993.
- M. Gorfine, R. Heller, and Y. Heller. Comment on detecting novel associations in large data sets by Reshef et al., science. Available: <http://iew3technionacil/~gorfinm/files/science6.pdf>, 2011.
- A. Gretton and L. Györfi. Consistent nonparametric tests of independence. *Journal of Machine Learning Research*, 11:1391–1423, 2010.
- A. Gretton, K.M. Borgwardt, M.J. Rasch, B. Schölkopf, and A. Smola. A kernel method for the two-sample problem. *Advances in Neural Information Processing Systems (NIPS)*, 19, 2007.
- A. Gretton, K. Fukumizu, C.H. TEO, L. Song, B. Schölkopf, and A. Smola. A kernel statistical test of independence. *Advances in Neural Information Processing Systems*, 20:585–592, 2008.
- A. Gretton, K.M. Borgwardt, M.J. Rasch, Schölkopf, and A. Smola. A kernel two-sample test. *The Journal of Machine Learning Research*, 13:723–773, 2012a.
- A. Gretton, D. Sejdinovic, H. Strathmann, S. Balakrishnan, M. Pontil, K. Fukumizu, and B.K. Sriperumbudur. Optimal kernel choice for large-scale two-sample tests. In *Advances in neural information processing systems*, pages 1205–1213, 2012b.

- Z Harchaoui, F. Bach, and E. Moulines. Testing for homogeneity with kernel fisher discriminant analysis. *Advances in Neural Information Processing Systems (NIPS), long version: arXiv:0804.1026v1*, 20:609–616, 2008.
- R. Heller, Y. Heller, and M. Gorfine. A consistent multivariate test of association based on ranks of distances. *Biometrika*, 100(2):503–510, 2013.
- W. Hoeffding. A non-parametric test of independence. *The Annals of Mathematical Statistics*, 19(4):546–557, 1948a.
- W. Hoeffding. A non-parametric test of independence. *The Annals of Mathematical Statistics*, 19(4):546–557, 1948b.
- T.R. Hughes, M.J. Marton, A.R. Jones, C.J. Roberts, R. Stoughton, C.D. Armour, H.A. Bennett, E. Coffey, H. Dai, Y.D. He, et al. Functional discovery via a compendium of expression profiles. *Cell*, 102(1):109–126, 2000.
- B. Jiang, C. Ye, and J. Liu. Non-parametric k -sample tests via dynamic slicing. *Journal of the American Statistical Association*, DOI:10.1080/01621459.2014.920257, 2014.
- J. Kinney and G. Atwal. Equitability, mutual information, and the maximal information coefficient. *Proceedings of the national academy of sciences of the USA*, doi:10.1073, 2014.
- B. Laurent and P. Massart. Adaptive estimation of a quadratic functional by model selection. *The Annals of Statistics*, 28(5):1302–1338, 2000.
- E.L. Lehmann and J.P. Romano. *Testing Statistical Hypotheses, 3rd Edition*. Springer, New York, 2005.
- M. Newton. Introducing the discussion paper by Szekely and Rizzo. *The Annals of Applied Statistics*, 3(4):1233–1235, 2009.
- L. Paninski. Estimation of entropy and mutual information. *Neural Computation*, 15(6):1191–1253, 2003.
- B. Póczos, Z. Ghahramani, and J. Schneider. Copula-based kernel dependency measures. In *International Conference on Machine Learning (ICML)*, 2012.
- A.N. Pettitt. A two-sample Anderson–Darling rank statistic. *Biometrika*, 63(1):161–168, 1976.
- D.N. Reshef, Y.A. Reshef, H.K. Finucane, S.R. Grossman, G. McVean, P.J. Turnbaugh, E.S. Lander, M. Mitzenmacher, and P.C. Sabeti. Detecting novel associations in large data sets. *Science*, 334(6062):1518–1524, 2011.
- F.W. Scholz and M.A. Stephens. K -sample Anderson-Darling tests. *Journal of the American Statistical Association*, 82(399):918–924, 1987.
- D. Sejdinovic, B. Sriperumbudur, A. Gretton, and K. Fukumizu. Equivalence of distance-based and rkhs-based statistics in hypothesis testing. *Annals of Statistics*, 41(5):2263–2291, 2013.

- N. Simon and R. Tibshirani. Comment on detecting novel associations in large data sets by Reshef et al., science. *arXiv:1401.7645*, 2011.
- R. Steuer, J. Kurths, C. Daub, J. Weise, and J. Selbig. The mutual information: detecting and evaluating dependencies between variables. *Bioinformatics*, 18(suppl 2):S231–S240, 2002.
- G. Székely and M. Rizzo. Testing for equal distributions in high dimensions. *InterStat*, 2004.
- G. Székely and M. Rizzo. Brownian distance covariance. *The Annals of Applied Statistics*, 3 (4): 1236–1265, 2009a.
- G. Székely, M. Rizzo, and N. Bakirov. Measuring and testing dependence by correlation of distances. *The Annals of Statistics*, 35:2769–2794, 2007.
- Gabor J Székely and Maria L Rizzo. Brownian distance covariance. *The annals of applied statistics*, 3(4):1236–1265, 2009b.
- O. Thas and J. Ottoy. An extension of the Anderson-Darling k-sample test to arbitrary sample space partition sizes. *Journal of Statistical Computation and Simulation*, 74(9):651–665, 2007.
- O. Thas and J.P. Ottoy. A nonparametric test for independence based on sample space partitions. *Communications in Statistics - Simulation and Computation*, 33 (3):711–728, 2004.
- Q. Vu, B. Yu, and R. Kass. Coverage-adjusted entropy estimation. *Statistics in medicine*, 26: 4039–4060, 2007.
- K. Yu, F. Liang, J. Ciampa, and N. Chatterjee. Efficient p-value evaluation for resampling-based tests. *Biostatistics*, 12(3):582–593, 2011.