

A QUANTITATIVE HEPPE'S THEOREM AND MULTIVARIATE BERNOULLI DISTRIBUTIONS

RICARDO FRAIMAN, LEONARDO MORENO, AND THOMAS RANSFORD

ABSTRACT. Using some extensions of a theorem of Heppes on finitely supported discrete probability measures, we address the problems of classification and testing based on projections. In particular, when the support of the distributions is known in advance (as for instance for multivariate Bernoulli distributions), a single suitably chosen projection determines the distribution. Several applications of these results are considered.

1. INTRODUCTION

The problem of dimension reduction is of increasing importance in modern data analysis. In this setting, we address the problem of when we can efficiently perform some important statistical analysis on high-dimensional discrete data, such as learning or hypothesis testing.

For supervised classification problems with high-dimensional binary data, nearest-neighbour or kernel based rules have a poor behaviour, and some other well-known techniques like Random Forest (Breiman (2001)) or Support Vector Machines (SVM see Boser et al. (1992)) are alternatives used in practice. However, as mentioned in Xu et al. (2021), there are just a few procedures particularly well designed for discrete data (see Needell et al. (2018); Molitor et al. (2019); Molitor and Needell (2021)). In particular, Needell et al. (2018) consider a classification rule for binary data, where the classification rule is updated recursively using projections, but over hyperplanes determined by subsets of the original coordinates chosen at random.

In this article, we propose an alternative approach that provides competitive procedures for the two main problems under consideration, hypothesis testing and learning. It also has potential applications to other statistical problems, since it provides a general dimension-reduction technique, based on a bound for the total variation distance between probability measures in terms of the distances between certain of their lower-dimensional projections. For example, this idea can be applied to perform clustering or to use depth functions in lower dimensions, among other applications that are not considered in the present manuscript.

2020 *Mathematics Subject Classification.* 60E05, 62E10, 62G10, 62H30, 62H15.

Key words and phrases. Classification, Discrete tomography, Heppes theorem, Multivariate Bernoulli, Random projections, Testing hypothesis.

Our approach is based on a variant of a theorem of Cramér and Wold (Cramér and Wold (1936)). According to the original Cramér–Wold theorem, if P and Q are Borel probability measures on \mathbb{R}^d whose projections P_L, Q_L onto each line L satisfy $P_L = Q_L$, then $P = Q$. There are several extensions of this theorem, in which one assumes more about the nature of the measures P, Q and less about the set of lines L for which $P_L = Q_L$. For example, if P and Q have moment generating functions that are finite in a neighbourhood of the origin, and if $P_L = Q_L$ for all lines L in a set of positive measure (in a natural sense), then $P = Q$. Articles on this subject include Rényi (1952), Gilbert (1955), Bélisle et al. (1997) and Cuesta-Albertos et al. (2007).

If P and Q are supported on finite sets, then it is even possible to differentiate between them using only a finite set of projections. The following result is due to Heppes (Heppes, 1956, Theorem 1'). We denote by P_H the projection of a probability measure P onto a general subspace H .

Theorem 1.1. *Let Q be a discrete probability measure on \mathbb{R}^d whose support consists of k points. Let H_1, H_2, \dots, H_{k+1} be subspaces of \mathbb{R}^d respectively of dimensions m_1, \dots, m_{k+1} , such that no two of these subspaces are contained in a single hyperplane, that is, no arbitrary straight line in \mathbb{R}^d can be perpendicular to more than one of the H_i . If P is a Borel probability measure on \mathbb{R}^d such that $P_{H_i} = Q_{H_i}$, $i = 1, \dots, k + 1$, then $P = Q$.*

An example of Rényi (1952) shows that a minimum of $(k + 1)$ subspaces are needed in Theorem 1.1. On the other hand, the ambient dimension d plays no role in this result, which could just as easily have been formulated in any separable Hilbert space. We also note that, though Q is assumed discrete, no such assumption is made about P . This makes the result suitable for hypothesis testing.

We establish a quantitative version of Heppes' theorem (see Theorem 2.1 in next section), where we show that the total variation distance between P and Q is at most the sum of the total variation distances between their projections. We also prove a refinement of Heppes' theorem where, if we know in advance the support of the distribution, then merely one suitably chosen projection suffices to determine the probability measure (see Proposition 2.3 in the next section). This projection just needs to be chosen to avoid a finite or countable number of bad directions, and thus almost all directions will be adequate. In particular, this is the case of multivariate binary distributions, which will be the focus of part of this paper.

The main results regarding projections are presented in Section 2. In Section 3 we propose a procedure for learning from projections, and adapt it for the learning problem for discrete tomography. Next we consider the case of testing for multivariate binary data in Section 4 and analyze the testing problem for some well-known distributions for multivariate binary data, an important problem for many applications (see for instance the nice review on this problem in Balakrishnan and Wasserman (2018)). Section 5 addresses a

different problem, namely if we can still say something when we have only one realization \mathbf{X} of a multivariate Bernoulli distribution.

Our results are illustrated with some simulation examples in Section 6 and three real-data examples in Section 7.

We end this section with a remark about possible extensions of our ideas to continuous distributions. As mentioned above, the starting point of this article is Heppes' theorem, which allows us to differentiate between two finitely supported measures on \mathbb{R}^d using only finitely many projections. For more general measures, it is usually necessary to use an infinite set of projections (see e.g. Hamedani and Tata (1975)). However, it turns out that finitely many projections suffice if both P and Q are elliptical distributions (a class of measures that includes not only Gaussian distributions but also many other continuous multivariate distributions of interest, including some with infinite first moments). This result, which may be viewed as a kind of analogue of Heppes' theorem for continuous distributions, is discussed in more detail at the end of the paper.

2. SOME VARIANTS OF HEPPES' THEOREM

Given two probability measures P, Q on a measurable space (E, \mathcal{A}) , we consider the *total variation distance* between them, namely

$$d_{TV}(P, Q) := \sup_{A \in \mathcal{A}} |P(A) - Q(A)|.$$

Also, given a Borel probability measure P on \mathbb{R}^d and a subspace H of \mathbb{R}^d , we write P_H for the projection of P onto H , namely the Borel probability measure on H given by

$$P_H(B) := P(\pi_H^{-1}(B)),$$

where $\pi_H : \mathbb{R}^d \rightarrow H$ is the orthogonal projection of \mathbb{R}^d onto H .

Clearly, if P, Q are Borel probability measures on \mathbb{R}^d and H is a subspace of \mathbb{R}^d , then

$$d_{TV}(P_H, Q_H) \leq d_{TV}(P, Q).$$

In this section we investigate some inequalities going in the opposite sense.

2.1. A quantitative Heppes theorem.

The following result is a quantitative generalization of Theorem 1.1.

Theorem 2.1. *Let Q be a discrete probability measure on \mathbb{R}^d whose support contains at most k points. Let H_1, \dots, H_{k+1} be subspaces of \mathbb{R}^d such that $H_i^\perp \cap H_j^\perp = \{0\}$ whenever $i \neq j$. Then, for every Borel probability measure P on \mathbb{R}^d , we have*

$$(1) \quad d_{TV}(P, Q) \leq \sum_{j=1}^{k+1} d_{TV}(P_{H_j}, Q_{H_j}).$$

The main idea of the proof is embodied in the following lemma, which is a useful result in its own right.

Lemma 2.2. *Let P, Q and H_1, \dots, H_{k+1} be as in the statement of the theorem. Then*

$$(2) \quad P(\{x\}) - Q(\{x\}) \leq \max_{1 \leq j \leq k+1} \left(P_{H_j}(\{\pi_{H_j}(x)\}) - Q_{H_j}(\{\pi_{H_j}(x)\}) \right) \quad (x \in \mathbb{R}^d).$$

Proof. Let $x \in \mathbb{R}^d$. Since the $(k+1)$ sets $(x + H_j^\perp) \setminus \{x\}$, $(j = 1, 2, \dots, k+1)$ are pairwise disjoint, and since Q is supported on a set containing at most k points, there exists a j such that $(x + H_j^\perp) \setminus \{x\}$ is disjoint from the support of Q . Hence

$$Q(x + H_j^\perp) = Q(\{x\}).$$

We then have

$$\begin{aligned} P(\{x\}) - Q(\{x\}) &= P(\{x\}) - Q(x + H_j^\perp) \\ &\leq P(x + H_j^\perp) - Q(x + H_j^\perp) \\ &= P_{H_j}(\{\pi_{H_j}(x)\}) - Q_{H_j}(\{\pi_{H_j}(x)\}). \end{aligned}$$

This establishes (2). □

Proof of Theorem 2.1. For each $j \in \{1, 2, \dots, k+1\}$, let A_j be the set of $x \in \mathbb{R}^d$ such that $(x + H_j^\perp) \setminus \{x\}$ is disjoint from the support of Q . The proof of Lemma 2.2 shows that $\cup_{j=1}^{k+1} A_j = \mathbb{R}^d$. Also, if B_j is a Borel subset of A_j , then $(B_j + H_j^\perp) \setminus B_j$ is contained in $\cup_{x \in A_j} ((x + H_j^\perp) \setminus \{x\})$, which is disjoint from the support of Q , and hence

$$Q(B_j + H_j^\perp) = Q(B_j).$$

Now let P be a Borel probability measure on \mathbb{R}^d , and let B be an arbitrary Borel subset of \mathbb{R}^d . Set $B_1 := B \cap A_1$ and $B_j := B \cap (A_j \setminus A_{j-1})$ for $j = 2, \dots, k+1$. Then B_1, \dots, B_{k+1} is a Borel partition of B such that $B_j \subset A_j$ for all j . Hence

$$\begin{aligned} P(B) - Q(B) &= \sum_{j=1}^{k+1} \left(P(B_j) - Q(B_j) \right) = \sum_{j=1}^{k+1} \left(P(B_j) - Q(B_j + H_j^\perp) \right) \\ &\leq \sum_{j=1}^{k+1} \left(P(B_j + H_j^\perp) - Q(B_j + H_j^\perp) \right) \\ &= \sum_{j=1}^{k+1} \left(P_{H_j}(\pi_{H_j}(B_j)) - Q_{H_j}(\pi_{H_j}(B_j)) \right) \\ &\leq \sum_{j=1}^{k+1} d_{TV}(P_{H_j}, Q_{H_j}). \end{aligned}$$

Also, since P, Q are both probability measures, we have

$$Q(B) - P(B) = P(\mathbb{R}^d \setminus B) - Q(\mathbb{R}^d \setminus B) \leq \sum_{j=1}^{k+1} d_{TV}(P_{H_j}, Q_{H_j}).$$

Finally, combining the last two inequalities, we obtain (1). \square

2.2. A Heppes-type result for distributions with pre-determined support.

If we know in advance the location of the supports of P and Q , rather than merely their cardinality, then it suffices to use just one judiciously chosen subspace H . This is even true if P, Q have countable supports, as the following result shows.

Proposition 2.3. *Let P, Q be probability measures on \mathbb{R}^d supported on a countable set $E \subset \mathbb{R}^d$. Let H be a subspace of \mathbb{R}^d such that $H^\perp \cap (E - E) = \{0\}$. Then*

$$d_{TV}(P, Q) = d_{TV}(P_H, Q_H).$$

In particular, $P_H = Q_H$ implies $P = Q$.

Remark. We can even choose H to have dimension equal to 1. Indeed, this amounts to ensuring that the hyperplane H^\perp misses the countable set $E - E$.

Lemma 2.4. *Let E be a countable subset of \mathbb{R}^d and let H be a subspace of \mathbb{R}^d with the property that $H^\perp \cap (E - E) = \{0\}$. Then the restriction of π_H to E is injective.*

Proof. If $x, y \in E$ and $\pi_H(x) = \pi_H(y)$, then $x - y \in (E - E) \cap H^\perp = \{0\}$, so $x = y$. \square

Proof of Proposition 2.3. Let B be a Borel subset of \mathbb{R}^d . From the lemma, we have $E \cap B = E \cap \pi_H^{-1}\pi_H(E \cap B)$, so

$$\begin{aligned} P(B) &= P(E \cap B) = P(E \cap \pi_H^{-1}\pi_H(E \cap B)) \\ &= P(\pi_H^{-1}\pi_H(E \cap B)) = P_H(\pi_H(E \cap B)), \end{aligned}$$

and likewise for Q . Hence

$$|P(B) - Q(B)| = |P_H(\pi_H(E \cap B)) - Q_H(\pi_H(E \cap B))| \leq d_{TV}(P_H, Q_H).$$

It follows that $d_{TV}(P, Q) \leq d_{TV}(P_H, Q_H)$. The reverse inequality is obvious, so we have equality. \square

A simple illustration of Proposition 2.3. Suppose that we have a random vector $(X_1, X_2) \in \mathbb{R}^2$, with $X_1 \sim \text{Ber}(p_1)$ and $X_2 \sim \text{Ber}(p_2)$. In this case, $E = \{(0, 0), (1, 0), (0, 1), (1, 1)\}$. We consider one-dimensional subspaces H , and we look at the orthogonal projections of the points in the support E on H , see Figure 1. There are only 4 subspaces (in red in Figure 1) that do not fulfill the condition $H^\perp \cap (E - E) = \{0\}$, namely those given by $\{(x, y) \in \mathbb{R}^2, ax + by = 0\}$ with $(a, b) \in \{(1, 0), (1, 1), (0, 1), (1, -1)\}$.




FIGURE 1. Illustration of Proposition 2.3.

2.3. Relation with other metrics.

Theorem 2.1 and Proposition 2.3 are expressed in terms of the total variation metric d_{TV} . For discrete distributions P, Q on \mathbb{R}^d , this metric is also given by

$$\begin{aligned} d_{TV}(P, Q) &= \frac{1}{2} \sum_x |P(\{x\}) - Q(\{x\})| \\ &= \inf \left\{ \mathbb{P}(X \neq Y) : (X, Y) \text{ such that } X \sim P \text{ and } Y \sim Q \right\}. \end{aligned}$$

There are several other commonly used metrics on the space of probability measures; see Gibbs and Su (2002) for an account of these. We mention one that will be used in what follows.

The *Wasserstein–Kantorovich metric* of order $p \geq 1$, defined for distributions on \mathbb{R}^d belonging to L^p , is given by

$$d_{W,p}(P, Q) := \inf \left\{ \mathbb{E}(|X - Y|^p)^{1/p} : (X, Y) \text{ such that } X \sim P \text{ and } Y \sim Q \right\}.$$

For one-dimensional distributions, an equivalent formulation is given by

$$(3) \quad d_{W,p}(P, Q) = \left(\int_0^1 |F^{-1}(t) - G^{-1}(t)|^p \right)^{1/p},$$

where F and G stands for the distribution functions of P and Q .

From (Gibbs and Su, 2002, Theorem 4), if P, Q are both supported on a bounded set $E \subset \mathbb{R}^d$ of diameter $\text{diam}(E)$, then

$$d_{W,1}(P, Q) \leq \text{diam}(E) d_{TV}(P, Q),$$

and, if E is finite and $d_{\min}(E)$ denotes the minimum separation between points of E , then

$$(4) \quad d_{W,1}(P, Q) \geq d_{\min}(E) d_{TV}(P, Q).$$

It follows that, if (P_n) is a sequence of probability distributions all supported on the same finite set, then P_n converges to P weakly if and only if $d_{TV}(P_n, P) \rightarrow 0$.

For the case of discrete tomography problem we will consider Mallow's L^2 -distance between histograms Mallows (1972), based on the Wasserstein distance given in (3). For histograms, the article Irpino et al. (2006) provides an exact calculation of this distance by an efficient algorithm.

3. CLASSIFICATION BASED ON PROJECTIONS AND DISCRETE TOMOGRAPHY

Discrete tomography, a term introduced by Larry Shepp in 1994, focuses on the reconstruction of binary (black and white) images based on a finite set of projections of the data, see for instance Gardner et al. (1996).

Let $E \subset \mathbb{Z}^d$ be the domain of the binary images, let $N := \text{card}(E)$, and let F be a subset of E . In this setting we only have access to the information given on parallel lines orthogonal to certain directions u of the images. More precisely, given a direction u in the unit sphere, the X -ray of the set F in direction u provides the number of points in the set on lines parallel to u , that is, if L_u stands for the line through the origin parallel to u , we are given the function

$$(5) \quad T_u(F)(x) := \text{card}(F \cap (x + L_u)), \quad x \in \mathbb{R}.$$

Given a direction u , we estimate $T_u(F)(x)$ based on an histogram built up from a partition of \mathbb{R} B_1, \dots, B_m as follows: for $x \in B_j$, define

$$\widehat{T}_u(F)(x) := \text{card} \left\{ \bigcup_{y \in B_j} (F \cap (y + L_u)) \right\}.$$

So we have the histogram of the points determined by the binary image of the points in the direction u , see Figure 2.




FIGURE 2. Representation of the histogram of the projection of the data in the direction u .

3.1. An algorithm with random projections. First we consider the following general classification problem into m classes: given an iid training sample, $\mathcal{D}^n = (\mathbf{X}_n, \mathbf{Y}_n) = \{(X_1, Y_1), \dots, (X_n, Y_n)\}$ with the same distribution as $(X, Y) \in E \times \{1, \dots, m\}$, we want to classify a new data X_{n+1} for which the label Y_{n+1} is unknown. Let P_l , $l = 1, \dots, m$ stand for the unknown

discrete distributions of $X|Y = l$ and let P_{nl} be the respective empirical distributions based on the training sample.

Let $\eta(x)$ denote the conditional mean of Y given $X = x$; namely, $\eta(x) = \mathbb{E}(Y|X = x)$. It is well known that the optimal rule for classifying a single new datum X is given by the Bayes' rule, which, for binary classification, reduces to $g^*(X) = \mathbb{I}_{\{\eta(X) \geq 1/2\}}$. Bayes' rule for finitely supported distributions to classify x corresponds to defining

$$(6) \quad \begin{aligned} g^*(x) = \tilde{Y} &:= \arg \max_{l=1, \dots, m} \frac{P(X = x|Y = l)P(Y = l)}{P(X = x)} \\ &= \arg \max_{l=1, \dots, m} \frac{P_l(X = x)P(Y = l)}{P_n(X = x)}, \end{aligned}$$

and a plug-in rule is given by

$$(7) \quad g_n(X_{n+1}, \mathcal{D}_n) = \hat{Y}_{n+1} := \arg \max_{l=1, \dots, m} \frac{P_{nl}(X = x)P_n(Y = l)}{P_n(X = x)},$$

which, for the binary classification case, reduces to

$$(8) \quad \begin{aligned} g^*(x) &= \mathcal{I}_{\{P_1(X=x)P(Y=1) > P_0(X=x)P(Y=0)\}}; \\ g_n(x) &= \mathcal{I}_{\{P_{n1}(X=x)P_n(Y=1) > P_{n0}(X=x)P_n(Y=0)\}}. \end{aligned}$$

Since we only have access to the projections of the distributions, our first approach is to make use of the relation (2). Let \mathcal{D}_j^{nl} , $l = 1, \dots, m$, $j = 1, \dots, k + 1$ be the projections of the training sample with label l on the subspace H_j , and let $X_{n+1,j}$ be the projection of a new data on H_j . To simplify the notation, we will consider the binary classification problem, $m = 2$.

A possible classification procedure will be to find under which member of the family of empirical distributions P_{nj} , $j = 1, \dots, m$ the new data X_{n+1} "fits better". For that purpose, we consider the distributions $\tilde{P}_{j(n+1)}$, $j = 1, \dots, m$ defined by just adding the data $X_{n+1} = x$ to each one of the j empirical distributions. Next we consider the vector of distances $d_{TV}(P_{nj}, \tilde{P}_{(n+1)j})$, $j = 1, \dots, m$ and we classify the point X_{n+1} into the class j for which the distance is minimal. We will use this last approach in what follows. Consistency is shown in the next theorem, where, for notational simplicity, we consider the binary classification problem, where now $\mathbb{E}(Y|X = x) = P(Y = 1|X = x)$. Let $L^* = P(g^*(X) \neq Y)$ be the Bayes risk, i.e. the risk of the best classifier, and let $L_n = L(g_n) = P(g_n(X, \mathcal{D}_n) \neq Y|\mathcal{D}_n)$. Recall that the classifier is weakly consistent if

$$\mathbb{E}(L_n) = P(g_n(X, \mathcal{D}_n) \neq Y) \rightarrow L^*, \text{ as } n \rightarrow \infty,$$

and strongly consistent if

$$\lim_{n \rightarrow \infty} L_n = L^* \text{ with probability 1.}$$

See for instance Devroye et al. (1996) for further details. Since (4) implies that $d_{TV}(P_{n0}, P_0)$ and $d_{TV}(P_{n1}, P_1)$ converge to zero a.s., it is clear that the plug-in classification rule given by equation (8) is strongly consistent.

Next we will prove the consistency of the new proposal rule given by:

$$(9) \quad g_n(x_{n+1}) := \begin{cases} 1 - i, & i = 0, 1, & \text{if } P_{ni}(x_{n+1}) = 0, \\ \mathcal{I}_{\{d_{TV}(P_{n1}, \tilde{P}_{(n+1)1}) < d_{TV}(P_{n0}, \tilde{P}_{(n+1)0})\}}, & \text{otherwise.} \end{cases}$$

Observe that the new rule (9) does not require us to have any notion of distance in the space of covariates E .

Theorem 3.1. *Let $\mathcal{D}^n = (\mathbf{X}_n, \mathbf{Y}_n) = \{(X_1, Y_1), \dots, (X_n, Y_n)\}$ be an iid training sample with the same distribution as $(X, Y) \in E \times \{0, 1\}$. Then the classification rule given by equation (9) is strongly consistent provided that the training sample satisfies the condition that $Y(x) = 1$ for all x such that $P_1(x) > P_0(x)$.*

Proof. First observe that since the support is finite, if $\tilde{P}_{n0}(x_{n+1}) = 0$ for n large enough, then $P_0(x_{n+1}) = 0$, so we can replace P_{ni} , $i = 0, 1$ by P_i in the top line of (9). Next observe that it suffices to show that, if

$$(10) \quad \mathcal{I}_{\{P_1(X=x) > P_0(X=x)\}} = 1,$$

then $\lim_{n \rightarrow \infty} \mathcal{I}_{\{d_{TV}(P_{n1}, \tilde{P}_{(n+1)1}) < d_{TV}(P_{n0}, \tilde{P}_{(n+1)0})\}} = 1$ with probability 1.

Since $d_{TV}(P_{n0}, P_0)$ and $d_{TV}(P_{n1}, P_1)$ converge to 0 a.s., the problem reduces to showing that equation (10) implies that

$$\lim_{n \rightarrow \infty} \mathcal{I}_{\{d_{TV}(P_1, \tilde{P}_{(n+1)1}) < d_{TV}(P_0, \tilde{P}_{(n+1)0})\}} = 1$$

with probability 1.

We only have to consider the case where $P_0(x_{n+1}) \neq 0$. We have $d_{TV}(P_1, \tilde{P}_{(n+1)1}) \rightarrow 0$ a.s., while, with probability one,

$$d_{TV}(P_0, \tilde{P}_{(n+1)0}) > \frac{1}{2} \left| P_0(x_{n+1}) - \frac{1}{(n_0 + 1)} \right| > P_0(x_{n+1})/4$$

for all n large enough. □

Since we only have access to the projections of the distributions, we make use of the bound based on projections given in the inequality (1) and replace the distances by their bounds.

Let \mathcal{D}_j^{nl} , $l = 1, \dots, m$, $j = 1, \dots, k + 1$ be the projections of the training sample with label l on the subspace H_j and let $X_{n+1,j}$ be the projection of a new datum on H_j .

Algorithm 1: Classification algorithm for learning from projections.

Input H_1, \dots, H_{k+1} ; \mathcal{D}_j^{nl} , $l = 1, \dots, m$, $j = 1, \dots, k + 1$ and
 $X_{n+1,j}$ $j = 1 \dots k + 1$;
 For each pair (j, l) , calculate the empirical distribution of \mathcal{D}_j^{nl} , $P_{\mathcal{D}_j^{nl}}$;
 Add the new projected data $X_{n+1,j}$ to each pair (j, l) of distributions
 $P_{\mathcal{D}_j^{nl}}$ and obtain $P_{\mathcal{D}_j^{(n+1)l}}$;
 Calculate $d_b(l) := \max_{j=1, \dots, k+1} d(P_{\mathcal{D}_j^{nl}}, P_{\mathcal{D}_j^{(n+1)l}})$, for $l = 1, \dots, m$,
 where d is a distance between the empirical distributions in \mathbb{R} ;
 Output $Y_{n+1} = l_s = \arg \min_{l=1, \dots, m} d_b(l)$.

Remark. In recent years, there has been a substantial literature on the subject of robustness for learning algorithms focussing on adversarial perturbations of the data (adversarial attacks), see for instance Bertsimas et al. (2019); López-Fidalgo and Wiens (2022), and the references therein. It treats the question of how changing a small fraction of data at arbitrary positions affects the learning strategies. Although distinct, this problem is related to the classical notion of breakdown point in robustness.

Among the kinds of perturbations considered, there are three main different settings: robustness against uncertainty in covariate features only in the testing sample, robustness against uncertainty only in labels, and robustness against uncertainty in covariate features and in the training sample, the last of these being the wildest. Although this problem is quite far from the central objective of the present article, it is not difficult to show that algorithm (9) is robust with respect to these kinds of perturbations. Indeed, from the proof of Theorem 3.1, one can derive that the procedure will be robust at x as long as the fraction α of adversarial perturbation points is smaller than $|P(x) - Q(x)|$. This condition seems to be mandatory for robustness for any classification rule from discrete data.

3.2. Discrete tomography: classification without reconstruction. In discrete tomography, for each F we only have access to the values $T_u(F)(x_i)$ for each direction u on a given grid $\mathbf{x} = \{x_1, \dots, x_M\}$. Once again, for notational simplicity, we consider the case of binary classification. Let $\mathcal{D}^n = \{\mathbb{F}_n, \mathbb{Y}_n\} = \{(F_1, Y_1), \dots, (F_n, Y_n)\}$ (the training sample) with the same distribution as $(F, Y) \in E \times \{0, 1\}$. Let $n_1 = \sum_{i=1}^n Y_i$ and $n_0 = n - n_1$ be the sizes of each class in the training sample. Denote by F_{ij} , $j = 1, \dots, n_i$, $i = 0, 1$, the data at each class according to its label i . We want to classify a new data F . Let $\mathbf{T}_u(F_{0j}) := (T_u(F_{0j})(x_1), \dots, T_u(F_{0j})(x_M))$, $j = 1, \dots, n_0$, and $\mathbf{T}_u(F_{1j}) := (T_u(F_{1j})(x_1), \dots, T_u(F_{1j})(x_M))$, $j = 1, \dots, n_1$, corresponding to direction u . Then define $\hat{\mathbf{T}}_u(F_{ij})$ as the histogram associated to the points of F_{ij} projected on the direction u . So, to each discrete tomography and direction, we associate the corresponding histogram.

We will denote by d the Mallow L^2 -distance between histograms. Next, based on this distance, for each direction u we apply the classical k -nearest

neighbour classification rule and then combine the results for the different directions u to produce the output. More precisely, we consider the following algorithm:

Algorithm 2: Classification algorithm for discrete tomography.

Input directions u_1, \dots, u_{k+1} ;

Input “projected” empirical histograms $\widehat{\mathbf{T}}_u(F_{ij})$ with

$j = 1, \dots, n_i, i = 0, 1$ for each direction u ;

Calculate the histogram corresponding to the new data F_{n+1} ;

Let d be a distance between histograms, given the direction u ,

determine $d\left(\widehat{\mathbf{T}}_u(F_{n+1}), \widehat{\mathbf{T}}_u(F_{ij})\right)$ for $j = 1, \dots, n_i, i = 0, 1$;

Fixed a number r of neighbours. By majority vote, the label

$Y_{n+1,u} = 0$ or 1 is assigned to the new observation by the k -nearest neighbour rule with $k = r$. This procedure is performed for each direction u ;

Output $Y_{n+1} = 0$ if $\frac{1}{k+1} \sum_{s=1}^{k+1} Y_{n+1,u_s} < 1/2$, and $Y_{n+1} = 1$ otherwise.

3.3. On the number of projections required in practice. If the support of the distributions is not known, but only its cardinality k , then both Heppes’ theorem and our Theorem 2.1 require the use of $(k + 1)$ projections. Although this value does not depend on the data dimension d directly, in practice, for high-dimensional binary data, the number of projections may be very large and therefore computationally hard. According to our experience, if one takes a sequence of increasing smaller values, the accuracy of the performance on the testing sample increases up to some point and then stabilizes. This suggest using a penalized procedure to select the number of random projections, providing a trade-off between the number of projections and the accuracy of the procedure.

4. MULTIVARIATE BINARY DATA

4.1. Notation. Let $\mathbf{X} = (X_1, \dots, X_d)$ be a random vector such that each coordinate $X_i, i = 1, \dots, d$ is a Bernoulli random variable, without any independence assumption about the coordinates X_i . The distribution of \mathbf{X} is called a multivariate Bernoulli distribution and has been studied by several authors, (see for instance Teugels (1990), Fontana and Semeraro (2018), Marchetti and Wermuth (2016), Euán and Sun (2020) Dai et al. (2013), and the references therein), where nice characterizations of the distribution are given. Also, the simulation problem for this type of distribution for different correlation structures is addressed in the literature, see for instance Oman (2009), Huber and Marić (2019), Jiang et al. (2021) and Fontana and Semeraro (2021).

The random vector \mathbf{X} takes values in the space

$$(11) \quad \Omega_d := \{0, 1\}^d = \{\omega = (y_1, \dots, y_d), y_i \in \{0, 1\}, i = 1, \dots, d\},$$

i.e. the set of d -tuples of zeros and ones, where y_i stands for the result of the i -th trial.

Let Δ_d be the set of all possible probabilities on Ω_d , i.e.

$$\Delta_d := \left\{ (p_1, \dots, p_{2^d}) : p_i \geq 0, \sum_{i=1}^{2^d} p_i = 1 \right\}.$$

For each probability $p \in \Delta_d$, given by

$$\{p(\omega) : \omega \in \Omega_d\},$$

we can derive the law of each X_i , as well as the law of $S_d := X_1 + \dots + X_d$, which are given by

$$q_i := P(X_i = 1) = \sum_{\{\omega \in \Omega_d : y_i = 1\}} p(\omega),$$

$$P(S_d = k) = \sum_{\{\omega \in \Omega_d : \sum_{i=1}^d y_i = k\}} p(\omega),$$

respectively.

Some nice and more explicit formulae are given in Dai et al. (2013), which can be used to perform a statistical analysis of the data in the case where we have a sample of iid vectors $\mathbf{X}_1, \dots, \mathbf{X}_N$ with a multivariate Bernoulli distribution.

4.2. Testing when we have an iid sample of high dimensional binary data. Suppose now that we are given a random sample $\{\mathbf{X}_1, \dots, \mathbf{X}_N\}$, where $\mathbf{X}_i = (X_{i1}, \dots, X_{id})$, for $i = 1, \dots, N$, and we want to perform some testing problems regarding this sample. This problem becomes difficult for high-dimensional data. However, several important problems fall into this category. This is the case, for instance, if we have a set of individuals N and for each of them X_{ij} represents the result of a given medical test (positive or negative), image analysis, or in marketing studies, among many others.

In this setting, since the support of the distributions is known in advance, we can make use of Proposition 2.3, which reduces the problem to checking a single projection on a appropriate subspace H . This will be particularly important in the case where we only have access to the projections of the data, as considered in the next sections.

First we choose a one-dimensional subspace fulfilling the condition in Proposition 2.3.

One-sample problem: Given an arbitrary model P_0 , and a sample $\{\mathbf{X}_1, \dots, \mathbf{X}_N\}$, where $\mathbf{X}_i = (X_{i1}, \dots, X_{id})$, for $i = 1, \dots, N$, let P_{0H} be the projection of P_0 on H , and let P_{NH} be the empirical distribution of the projections of the sample on H .

Then, the problem of testing $(P = P_0)$ vs $(P \neq P_0)$, reduces to testing $(P_H = P_{H0})$ vs $(P_H \neq P_{H0})$, which can be handled using

the Kolmogorov–Smirnov test or the Cramer–von-Mises test, among other possibilities.

For instance, for the Kolmogorov–Smirnov test, let $KS(P_{NH}, P_{H0}) := \sup_x \|F_{NH}(x) - F_{H0}(x)\|$, the Kolmogorov distance between the one-dimensional distribution functions of P_{NH} and P_{H0} . Reject the null assumption at level α if $KS(P_{NH}, P_{H0}) > c_\alpha$.

The Kolmogorov–Smirnov test for discrete data has been analyzed by several authors. See for instance Conover (1972) and Dimitrova et al. (2020). In the latter article, a software package is available to calculate the exact critical value c_α for discrete distributions.

Two-sample problem: Given two samples $\{\mathbf{X}_1, \dots, \mathbf{X}_N\}$, $\{\mathbf{Y}_1, \dots, \mathbf{Y}_M\}$ with distributions P and Q respectively, we want to test $(P = Q)$ vs $(P \neq Q)$. Let P_{NH} be the empirical distribution of the projections of the first sample on H , and let P_{MH} be the empirical distribution of the projections of the second sample on H .

Again the problem reduces to testing if $(P_H = Q_H)$ vs $(P_H \neq Q_H)$, which can be handled using the two-sample Kolmogorov–Smirnov test or the Cramer–von-Mises test, among other possibilities.

In this case, we consider the Kolmogorov distance between the two empirical distribution functions, i.e., $KS(P_{NH}, P_{MH}) := \sup_x \|F_{NH}(x) - F_{MH}(x)\|$, and reject the null assumption at level α if $KS(P_{NH}, P_{MH}) > c_\alpha$.

Some practical examples, like testing for the Poisson–Binomial distribution, are considered in Section 6 below.

5. BERNOULLI SERIES DISTRIBUTIONS AND THEIR SUMS

5.1. Background.

The problem we want to analyze in this section is whether we can still say something when we have only one realization \mathbf{X} of a multivariate Bernoulli distribution. In particular, we are interested in characterizing the law of the sum S_d for different distributions $p \in \Delta_d$, in particular those for which the law of S_d is close to a Binomial distribution $B(d, 1/2)$ with parameters d and $1/2$.

Let $L_d(p)(k) = P(S_d = k)$. Given $\epsilon > 0$ let

$$(12) \quad \Delta(d; \epsilon) := \left\{ p \in \Delta_d : \max_{0 \leq k \leq d} |L_d(p)(k) - B(d, 1/2, k)| \leq \epsilon \right\},$$

the set of probabilities $p \in \Delta_d$ which are uniformly close to $B(d, 1/2)$.

Next we choose at random a probability measure $p \in \Delta_d$ according to the normalized Lebesgue measure μ on the simplex Δ_d , which is related to the notion of physical entropy.

In this setting, we have the following result of Chevallier (2010).

Theorem 5.1 ((Chevallier, 2010, Theorem 1)). *There exists a constant $A \leq 2$ such that*

$$(13) \quad \mu(\Delta(d; \epsilon)) \geq 1 - \frac{A\sqrt{d}}{\epsilon^2 2^{d-1}}$$

and

$$(14) \quad \mu\left(\{p \in \Delta_d : \sup_{I \subset \{1, \dots, d\}} |L_d(p)(I) - B(d, 1/2, I)| \leq \epsilon\}\right) \geq 1 - \frac{Ad^{5/2}}{\epsilon^2 2^{d-1}}.$$

Remarks.

- 1) Observe that, since we are choosing at random a probability measure $p \in \Delta_d$ according to the normalized Lebesgue measure on the simplex Δ_d , the mean value $\mathbb{E}(S_d) = d/2$.
- 2) Even for moderate d , the bound is very close to 1, and the results are uniform. For instance, if $d = 30$ and $\epsilon = 0.01$, the right hand side of (13) is greater than 0.99979, and for (14), if $d = 50$ and $\epsilon = 0.01$, the bound is greater than 0.9999. For most probability distributions $p \in \Delta(d, \epsilon)$, the distribution of S_d will be much the same, and so little information can be derived from its distribution. In other words, with high probability the distribution of S_d will be very close to $B(d, 1/2)$ if we choose at random the joint distribution p of the random variables X_1, \dots, X_d . This means that the class of probabilities p which are not close to the $B(d, 1/2)$ is very small.
- 3) However the set of “rare” distributions p contains some well-known structured distributions. For instance if the random variables X_i are iid Bernoulli with common parameter $q \neq 1/2$, the law of S_d is not close to the $B(d, 1/2)$. Moreover, if the random variables are independent but with different parameters,

$$p(w) = \prod_{\{y_i=1, 1 \leq i \leq d\}} q_i \prod_{\{y_i=0, 1 \leq i \leq d\}} (1 - q_i), \quad w \in \Omega_d,$$

then the law of S_d is a Poisson–Binomial distribution, given by

$$P(S_d = k) = \sum_{A \in \Lambda_k} \prod_{i \in A} q_i \prod_{j \in A^c} (1 - q_j),$$

where Λ_k is the family of all subsets of $\{0, 1, \dots, k\}$. This typically is also far from $B(d, 1/2)$. For instance, this will be the case if $\bar{q} := E(S_d)/d = \frac{1}{d} \sum_{i=1}^d q_i \neq 1/2$, where $q_i = P(X_i = 1)$.

What if we assume that $\mathbb{E}(S_d) = d/2$?

Even if $\mathbb{E}(S_d) = d/2$, from Theorem 1 in Ehm (1991) we have that the total variation distance $d_{TV}(\cdot, \cdot)$ between the Binomial($d, 1/2$) and the Poisson Binomial distribution is larger than

$$\frac{\sum_{j=1}^d (q_j - 1/2)^2}{31d},$$

which will not be small except for the case where all but a vanishing fraction α_d of the q_j are very close to $1/2$.

From a statistical point of view we have to be very careful when dealing with high-dimensional binary data, if the data are not independent. Indeed, in contrast to the case where we have a Binomial(d, q) distribution, where we have a sample of size d of iid Bernoulli(q) random variables, we have only one sample of the distribution p in a high-dimensional binary space. However, these results provide an excellent way to perform some statistical tests for important applications. Indeed, we can use these results for testing if the distribution of S_d has some structure, which in terms of our setting will correspond to be in the set of “rare” distributions.

5.2. Testing when we have only one high-dimensional binary datum.

Given a sample X_1, \dots, X_d of Bernoulli random variables, let $S_d = \sum_{i=1}^d X_i$ and let $0 < \alpha < 1$. We start by fixing ϵ such that

$$(15) \quad 1 - \frac{2\sqrt{d}}{\epsilon^2 2^{d-1}} \geq 1 - \alpha/2.$$

For $\alpha = 0.01$ and $\epsilon = 0.01$ this will hold for any $d \geq 26$, for $\epsilon = 0.005$ for any $d \geq 28$ and for $\epsilon = 0.001$ for any $d \geq 36$.

We consider as the null assumption that the distribution p is not rare, i.e. that it belongs to the set $\Delta(d, \epsilon)$, a set that satisfies

$$(16) \quad \mu(\Delta(d; \epsilon)) \geq 1 - \alpha/2,$$

and next we look if the observed value for S_d is very atypical for a $B(d, 1/2)$.

Because of the definition of $\Delta(d, \epsilon)$, a conservative test of level α can be performed, rejecting the null assumption if

$$(17) \quad S_d \notin [l, r],$$

where $P(B(d, 1/2) \notin [l, r]) \leq \alpha/2$.

Rejecting from the right ray or from the left ray will have different implications in some applications.

In case that we reject the null hypothesis, we would like to characterize in some way the rare distribution, but it will not be possible with just a single sample of the vector of binary data. In order to do so, we need to have an iid (or at least mixing) sample of the binary vectors.

5.3. Testing for rare distributions.

We proceed as follows. For each vector \mathbf{X}_i let $S_{id} = \sum_{j=1}^d X_{ij}$ the sum of the coordinates of the vector \mathbf{X}_i , $i = 1, \dots, N$. For each $k = 0, \dots, d$, we consider the empirical probability given by

$$(18) \quad P_N(k) := \frac{1}{N} \sum_{i=1}^N \mathcal{I}_{\{S_{id}=k\}}.$$

From the Hoeffding inequality (Bennett (1962); Hoeffding (1963)), we can derive that, for any $t > 0$,

$$(19) \quad P(|P_N(k) - P(S_{1d} = k)| > t) \leq 2 \exp(-2Nt^2),$$

and so, from the union bound,

$$(20) \quad P\left(\max_{0 \leq k \leq d} |P_N(k) - P(S_{1d} = k)| > t\right) \leq 2(d+1) \exp(-2Nt^2).$$

Now we are ready to derive a conservative level- α test for the null hypothesis

H0: the distribution p is not rare.

This can be obtained from inequalities (13) and (19) or (20), rejecting the null assumption when:

$$(21) \quad |P_N(k) - B(d; 1/2, k)| > a,$$

where a is chosen to satisfy the inequalities (22) and (23) below, and where k is the observed value. Indeed, we have that

$$\begin{aligned} & P\left(|P_N(k) - B(d; 1/2, k)| > a\right) \\ & \leq P\left(|P_N(k) - P(S_{1d} = k)| > a/2\right) + P\left(|P(S_{1d} = k) - B(d; 1/2, k)| > a/2\right). \end{aligned}$$

The first term on the right hand side will be smaller than $\alpha/2$ provided that N and a are chosen so that

$$(22) \quad \exp(-2Na^2) < \alpha/2,$$

while the second one will be smaller than $\alpha/2$ if d and a are chosen so that

$$(23) \quad \frac{8\sqrt{d}}{a^2 2^{d-1}} < \alpha/2.$$

An easy calculation leads to

$$(24) \quad a = \max \left\{ \sqrt{\frac{-\log(\alpha/2)}{2N}}, \frac{d^{1/4}}{2^{(d-5)/2} \sqrt{\alpha}} \right\}.$$

For instance, if $d = 20$, $N = 200$ and $\alpha = 0.05$, then $a = \max(0.052, 0.086)$. For $d \geq 30$ the second term is negligible (< 0.0018), and the bound is driven by the first term.

6. SIMULATIONS

6.1. Example 1: Binary data classification. We generate two different multivariate Bernoulli distributions in dimensions $d \in \{5, 10, 15, 20\}$. Both distributions have marginals Bernoulli(1/2). In one of them, the components are independent, while, for the other one, we consider correlated components with a parameter Cor taking values in the set $\{0.1, 0.3, 0.5, 0.7, 0.9\}$ for each different scenario.

We generate 200 observations for each class and 25% of them are used for the testing sample. The data are generated as indicated in Park et al. (1996).

TABLE 1. Average misclassification errors for each scenario in the Example 6.1 over 1000 replicates. Standard errors are reported in brackets.

mean (sd)	RF				N-S-W				RP ₁₀₀₀₀			
	Dimension											
Corr.	5	10	15	20	5	10	15	20	5	10	15	20
0.1	46.7 (5.0)	46.9 (5.1)	46.5 (5.2)	46.1 (5.2)	46.2 (5.1)	47.2 (5.5)	47.7 (5.0)	46.8 (5.0)	46.6 (4.6)	47.4 (5.0)	46.5 (5.1)	45.8 (4.6)
0.3	34.8 (4.9)	32.7 (4.4)	32.0 (4.3)	31.7 (4.2)	36.0 (5.3)	33.3 (4.2)	32.6 (4.4)	31.5 (4.1)	35.4 (4.8)	34.0 (4.5)	31.7 (5.3)	31.8 (4.5)
0.5	24.5 (4.7)	19.6 (3.9)	18.8 (3.7)	17.9 (3.2)	24.7 (4.7)	20.7 (3.9)	19.2 (4.2)	18.3 (4.6)	24.8 (5.1)	20.2 (4.0)	19.3 (3.7)	17.6 (3.7)
0.7	17.1 (3.7)	10.1 (2.9)	7.4 (2.7)	6.9 (1.9)	17.5 (3.7)	10.8 (3.2)	9.1 (2.9)	8.5 (3.0)	17.2 (3.9)	10.7 (2.8)	9.1 (3.2)	8.5 (2.8)
0.9	9.7 (3.0)	3.9 (2.1)	3.3 (1.5)	3.0 (1.2)	9.7 (3.1)	4.5 (2.5)	2.8 (1.8)	2.5 (1.6)	9.7 (3.0)	3.9 (2.0)	2.8 (1.8)	2.2 (1.6)

The distances between the empirical distributions P_{n_i, i, u_i} and $P_{n_{i+1}, i, u}$ are calculated according to the statistic DTS, given for instance in Dowd (2020).

Next we apply Algorithm 1 to perform the classification. In all cases we use 100 one-dimensional projections. For each of the considered scenarios, 100 replicates are performed. The misclassification error of Random Forest, Needell et al. (2018) (denoted as N-S-W method) considering $L = 3$ “levels”, and our proposal are given in Table 1, while the boxplots of the misclassification errors of our proposal are reported in Figure 3 for different numbers of projections. One can see that, when the correlations of the second group become higher, the distributions are “further apart”, and the misclassification error becomes smaller. Also, we observe that, when the dimension of the space is higher, the algorithm performs better.




FIGURE 3. Boxplots of the misclassification errors for each scenario based on 100 replications.

6.2. Example 2: Classification of discrete Tomography by using distances between histograms. In this example we generate a set of “phantom images” using two different random patterns. In the first scenario, the first sample (of size 200) each image is built with 5 circles of centers

$(2, 2)$, $(-2, 2)$, $(2, -2)$, $(-2, -2)$ and $(0, 0)$ with a normal random radius with mean 1 and standard deviation $1/10$. The second sample (of size 200) is built up in the same way, but adding a circle centred at $(0, 2)$ and with a normal random radius with mean $1/2$ and standard deviation $1/10$ (see Figure 4). The points that define the image are in an equi-spaced grid at distance 0.05.




FIGURE 4. Representation of an image of each group in scenario 1.

For the second scenario, the first sample is equal to that of the first scenario. Now the second sample consists on 5 circles with centres $(2, 2)$, $(-2, 2)$, $(2, -2)$, $(-2, -2)$, $(0, 0)$ and random radius with mean 1.2 and standard deviation $1/10$, see Figure 5.

The value k for the k -NN is the one that minimizes the misclassification error in the testing sample ($k = 21$).

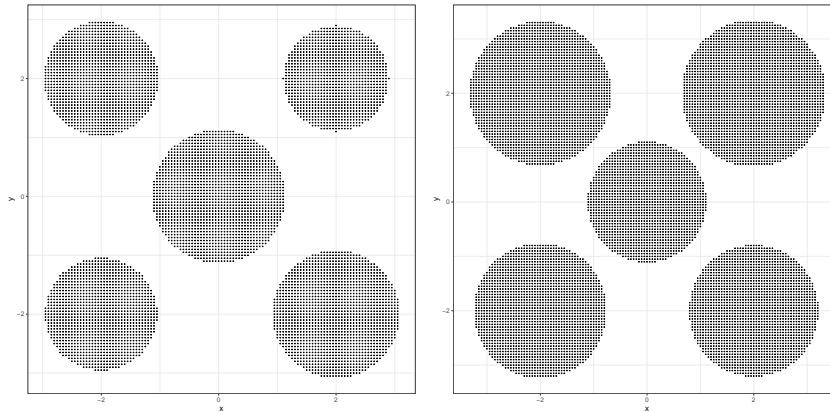


FIGURE 5. Representation of an image of each group in scenario 2.

75% of the samples are considered for the training sample and the rest for the testing sample. We use Algorithm 2 with 100 projections at random in each case.

For each direction we project the points of the image and build up an histogram. Given a testing sample using $k_N N$, we observe the proportion of neighbours at each group. If the average over all directions of the proportions of each group is greater for group 1, then the image is classified as being from group 1, otherwise is classified as from group 2. As mentioned above, the distance between the histograms is calculated as indicated in Irpino et al. (2006), i.e., using the Mallow distance in L^2 .

The misclassification rate in the testing sample are 2.55% and 6% in scenarios 1 and 2 respectively.

6.3. Example 3: A test for binary data. We consider a sample of Bernoulli vectors just as in section 4.2, $\{\mathbf{X}_1, \dots, \mathbf{X}_{200}\}$ where $\mathbf{X}_i = (X_{i1}, \dots, X_{id})$, for $i = 1, \dots, 200$, with $d = 8$, and all marginals $X_j \sim \text{Bernoulli}(1/2)$ for all j .

The testing problem is the following: under the null assumption, the components of the vector are independent, i.e.

$$P_0(X_1 = j_1, \dots, X_d = j_d) = \left(\frac{1}{2}\right)^d, \quad j_1, \dots, j_d \in \{0, 1\}.$$

Under the alternative the data are simulated using the odds ratio

$$OR_{ij} = \frac{P(X_i = 1, X_j = 1) P(X_i = 0, X_j = 0)}{P(X_i = 0, X_j = 1) P(X_i = 1, X_j = 0)} = \gamma,$$

with $\gamma \in [1, 3]$ if $i \neq j$ and $+\infty$ if $i = j$. The larger the value of γ , the further the alternative is from the null.

The data are projected into $k \in \{1, 10, 50, 100, 500\}$ random directions u_1, \dots, u_k . The testing statistic is given by $KS(\gamma) = \frac{1}{k} \sum_{j=1}^k KS(P_{Nu_j}, P_{H0})$. The distribution of the statistic under the null is obtained by Monte Carlo. The power function is obtained using 1000 replicates as a function of γ , see Figure 6.

Observe that, with just 50 projections, for $\gamma \geq 1.75$ the power of the test is greater than 73%. The simulations were performed using the R package **mipfp**, see Barthélemy and Suesse (2018).

6.4. Example 4: Test for Poisson-Binomial. Next we consider the case of Poisson-Binomial S_d distributions with $d \in \{50, 100, 200, 500, 1000\}$, i.e. the sum of d independent Bernoulli random variables with different parameters. The d probability parameters of the distribution are generated from a Beta distribution with parameters $\gamma_1 \in [2, 4]$ and $\gamma_2 = 2$. If $\gamma_1 = \gamma_2$, then

$$\frac{1}{d} E(S_d) = \frac{1}{d} \sum_{i=1}^d q_i \approx \frac{\gamma_1}{\gamma_1 + \gamma_2} = \frac{1}{2},$$

for d large enough.




FIGURE 6. Estimated power function for $\gamma \in [1, 3]$ for different number of projections.

As indicated in Section 5.2, the power of the test is estimated via Monte Carlo for each value of d as a function of γ_1 , see Figure 7. The behaviour becomes better when increasing the dimension d , while the problem becomes easier by increasing γ_1 .




FIGURE 7. Power function for $\gamma_1 \in [2, 4]$.

7. REAL-DATA EXAMPLES

In this section three examples with real data are exhibited. In both cases the data are a set of binary vectors and Algorithm 1 is performed.

7.1. Example 5. The data describe the diagnostic of Cardiac Single Proton Emission Computed Tomography (SPECT) images. Each patient is classified into one of two categories: normal and not normal. The database has 267 images called SPECT corresponding to each patient, “that measures radioactive counts that represent perfusion in the LV muscle in a specified ROE” (regions of interest) resulting 22 regions of interest, see Kurgan et al. (2001).

The images are processed and transformed into a vector of dimension 44 of continuous variables. Later on, this pattern is processed again to obtain a binary vector of dimension 22 for each patient.

The sample is split into 80 observations for the training sample and 187 for the testing sample. The database is available from the UCI repository¹.

We apply Algorithm 1 to classify the data using different numbers of random projections ($k \in \{5, 10, 20, 50, 100, 500, 1000, 10000\}$).

The procedure is replicated 100 times at each scenario and the boxplots of the misclassification rates are given in Figure 8.




FIGURE 8. Boxplot of the misclassification rates obtained from the testing sample for different numbers of random projections.

The performance of the method that we propose (denoted by RP_k) is compared with two classical classification methods (Random Forest and SVM) and in particular with the proposed methodology for binary data in N-S-W (with $L = 3$). Table 2 shows the confusion matrices obtained in the test sample for each method. The misclassification rates obtained from the testing sample are 23%, 24.1%, 19.3% and 18.4% for Random Forest, SVM

¹<https://archive.ics.uci.edu/ml/datasets/spect+heart>

TABLE 2. Confusion matrices (in percentages) in the test samples of binary classifications. Here 0 is normal and 1 is not normal.

	Labels											
	Random Forest		SVM		N-S-W		RP ₁₀₀₀₀					
	0	1	0	1	0	1	0	1				
Pred.	0	5.9	2.1	0	6.4	1.6	0	5.3	2.7	0	5.1	2.9
	1	20.9	71.1	1	22.5	69.5	1	16.6	75.4	1	15.5	76.5

,N-S-W, and RP_k respectively. Therefore, in this example, methods designed for discrete data perform better than some general classification methods. In addition, the N-S-W and RP_k methods have similar behaviour.

7.2. Example 6. In Grisoni et al. (2019), consensus machine learning algorithms are used to predict the binding to the androgen receptor AR. We quote from Grisoni et al. (2019):

“The nuclear androgen receptor (AR) is one of the most relevant biological targets of Endocrine Disrupting Chemicals (EDCs), which produce adverse effects by interfering with hormonal regulation and endocrine system functioning. The nuclear androgen receptor (AR), whose gene is located on the X chromosome, is expressed in a wide range of tissues and plays a fundamental biological role in bone, muscle, prostate, adipose tissue, and the reproductive, cardiovascular, immune, neural, and hemopoietic systems. AR is one of the target receptors of the so-called Endocrine Disrupting Chemicals (EDCs), exogenous compounds able to disturb hormonal regulation and the endocrine system functioning, thereby producing adverse effects in humans and wildlife. EDCs can interact directly with a given nuclear receptor and perturbate or modulate downstream gene expression, but they can also have direct effects on genes and epigenetic impact. Disruption of AR-mediated processes can cause irreversible consequences in human health. For example, some chemicals, like pesticides (e.g., DDT), disrupt male reproductive development and function by inhibiting androgen-receptor-mediated events. Recently, machine learning (ML) and computer-aided techniques have shown to be useful for modeling nuclear receptor modulation of chemicals at different levels, such as drug discovery and design and testing prioritization campaigns.”

The dataset, which is available at the UCI repository², contains agglutinate (positive) and no negative agglutinate molecules. 150 of each type are taken as training sample and 45 of each type as testing sample. Each molecule is represented by 1024 binary molecular fingerprints, that is, a 1024 length sequence of 0s and 1s, see Grisoni et al. (2019); Piir et al. (2021).

The boxplots of the misclassification errors are given in Figure 9 for different numbers of random projections.

²<https://archive.ics.uci.edu/ml/datasets/QSAR+androgen+receptor>




FIGURE 9. Boxplots for the misclassification errors at the testing sample for different number of random projections.

Considering 1000 projections, the performance of the classification in each class is evaluated through three indices in the test sample:

$$\text{Sensitivity: } S_n = \frac{TP}{TP+FN} = 0.97$$

$$\text{Specificity: } S_p = \frac{TN}{TN+FP} = 0.72$$

$$\text{Non-Error Rate: } NER = \frac{S_n+S_p}{2} = 0.845,$$

where TP and TN stand for the misclassification errors in each class. These results correspond to the median values over 100 replicates. We observe that, with respect to the results obtained in Grisoni et al. (2019), we get a slightly smaller value for the specificity, but better results for the sensitivity and the non-error rate.

7.3. Example 7: The MNIST database. This example is similar to Example 2, but with real data. The base has a set of handwritten digits and is available in <http://yann.lecun.com/exdb/mnist/>. The digits have been size-normalized and centered in a fixed-size image. Each image is represented by a matrix of $28 \times 28 = 784$ pixels. Each pixel takes an integer value from 0 to 255 (grey scale). The labels form the components of a vector representing the digit shown in the image. In this work the image is dichotomized to 1–0 (black and white respectively). If the pixel is above 100 on the grey scale we associate 1 to it, and otherwise 0. As an illustration, between the sample of digits 3, 4 and 9 in a test sample, the classification is performed (as in the Example 2). In Figure 10, three dichotomized images of each of these digits are shown.

The training sample size is 758, 736, 757 and the test sample size is 252, 246, 252 for digits 3, 4 and 9 respectively. The value $k = 19$ (for the k-NN) and 100 random projections are considered. The confusion matrix of the binary classification in the test sample are reported in Table 3.

Therefore, the overall prediction errors are 5.4%, 3.2% and 7.6% in the test sample for the classifications of digit pairs 3-4, 3-9 and 4-9 respectively.




FIGURE 10. Representation of dichotomized sample images: three digits 3 (first row), three digits 4 (second row) and three digits 9 (third row).

TABLE 3. Confusion matrices in the test samples of binary classifications.

	Labels								
	3	4	3	4	9	3	4	9	
Prediction	3	248	23	3	243	7	4	230	17
	4	4	223	9	9	245	9	16	235

8. CONCLUSIONS

We provide some new statistical techniques to treat discrete high-dimensional data, based on some extensions of Heppes' theorem, in particular a quantitative version of Heppes' theorem that bounds total variation distance between two probability measures in the high-dimensional space by the sum of the total variation distances between their projections.

We show that, if we know in advance the supports of the distributions (which is the case, for instance, of multivariate Bernoulli distributions), then it suffices to consider just one projection, provided that the direction is chosen to exclude a certain finite or countable set of bad directions, which are well determined.

Using the preceding results, we develop a new procedure for learning based on projections, and we adapt it to the case of discrete tomography, where we only have access to information related to the projections in a certain family of directions.

We also address the problem of testing for multivariate binary data, in particular for some well-known models like the Poisson-Binomial model. Our approach can be used for many other different problems for high-dimensional discrete data, for example: the Dirichlet-multinomial distribution, the multivariate hypergeometric distribution, the multivariate Pólya–Eggenberger distributions, and the negative multinomial distribution, among others (see

Johnson et al. (1997)). We further analyze the case where we only have one realization of a multivariate Bernoulli distribution.

Lastly, we perform a small simulation study and analyze three real-data examples. The results obtained are quite encouraging.

We conclude this article with a brief discussion of the potential extension of these ideas to continuous distributions. As mentioned in the introduction, in order to distinguish between general continuous distributions, it is usually necessary to use an infinite set of projections (see e.g. Hamedani and Tata (1975)). However, there is at least one case where finitely many projections will suffice. A Borel probability measure P on \mathbb{R}^d is called *elliptical* if its characteristic function has the form

$$\phi_P(\xi) = e^{i\mu \cdot \xi} \psi(\xi^T \Sigma \xi) \quad (\xi \in \mathbb{R}^d),$$

where Σ is a real positive semi-definite $d \times d$ matrix, μ is a vector in \mathbb{R}^d , and $\psi : [0, \infty) \rightarrow \mathbb{C}$ is a continuous function. Examples of elliptical distributions include multivariate versions of Gaussian, Student, Cauchy, Bessel, uniform and logistic distributions. The following theorem is a summary of results established in Fraiman et al. (2022).

Theorem 8.1. *Given $d \geq 2$, there exists a set \mathcal{L} of $(d^2 + d)/2$ lines in \mathbb{R}^d with the following property: if P, Q are elliptical distributions on \mathbb{R}^d such that $P_L = Q_L$ for all $L \in \mathcal{L}$, then $P = Q$. No smaller set of lines will do.*

Theorem 8.1 leads to the derivation of statistical tests for equality of elliptical distributions. This program is carried out in Fraiman et al. (2022). An interesting avenue for further exploration, suggested by one of the referees, would be to combine the two cases already studied (discrete and elliptical) to develop tests for distributions that are discrete distributions with weak noise, for example, vector sums of discrete distributions and elliptical distributions with small variance.

FUNDING STATEMENT

Fraiman and Moreno were supported by grant FCE-1-2019-1-156054, Agencia Nacional de Investigación e Innovación, Uruguay. Ransford was supported by grants from NSERC and the Canada Research Chairs program.

DATA AVAILABILITY STATEMENT

All the data that we use are publicly available; the precise sources are cited in the article. The source code files in R are available from the authors upon request.

REFERENCES

- Balakrishnan, S. and Wasserman, L. (2018) Hypothesis testing for high-dimensional multinomials: a selective review. *Ann. Appl. Stat.*, **12**, 727–749. URL: <https://doi-org.acces.bibl.ulaval.ca/10.1214/18-AOAS1155SF>.

- Barthélemy, J. and Suesse, T. (2018) mipfp: An R package for multidimensional array fitting and simulating multivariate Bernoulli distributions. *Journal of Statistical Software, Code Snippets*, **86**, 1–20.
- Bélisle, C., Massé, J.-C. and Ransford, T. (1997) When is a probability measure determined by infinitely many projections? *Ann. Probab.*, **25**, 767–786. URL: <https://doi-org.acces.bibl.ulaval.ca/10.1214/aop/1024404418>.
- Bennett, G. (1962) Probability inequalities for the sum of independent random variables. *J. Amer. Statist. Assoc.*, **57**, 33–45.
- Bertsimas, D., Dunn, J., Pawlowski, C. and Zhuo, Y. D. (2019) Robust classification. *INFORMS J. Optim.*, **1**, 2–34. URL: <https://doi-org.acces.bibl.ulaval.ca/10.1287/ijoo.2018.0001>.
- Boser, B. E., Guyon, I. M. and Vapnik, V. N. (1992) A training algorithm for optimal margin classifiers. In *Proceedings of the Fifth Annual Workshop on Computational Learning Theory, COLT '92*, 144–152. New York, NY, USA: Association for Computing Machinery. URL: <https://doi.org/10.1145/130385.130401>.
- Breiman, L. (2001) Random forests. *Machine Learning*, **45**, 5–32.
- Chevallier, N. (2010) Law of the sum of Bernoulli random variables. *Teor. Veroyatn. Primen.*, **55**, 87–103. URL: <https://doi-org.acces.bibl.ulaval.ca/10.1137/S0040585X97984644>.
- Conover, W. J. (1972) A Kolmogorov goodness-of-fit test for discontinuous distributions. *J. Amer. Statist. Assoc.*, **67**, 591–596. URL: [http://links.jstor.org.acces.bibl.ulaval.ca/sici?sici=0162-1459\(197209\)67:339<591:AKGTFD>2.0.CO;2-G&origin=MSN](http://links.jstor.org.acces.bibl.ulaval.ca/sici?sici=0162-1459(197209)67:339<591:AKGTFD>2.0.CO;2-G&origin=MSN).
- Cramér, H. and Wold, H. (1936) Some theorems on distribution functions. *J. London Math. Soc.*, **11**, 290–294. URL: <https://doi-org.acces.bibl.ulaval.ca/10.1112/jlms/s1-11.4.290>.
- Cuesta-Albertos, J. A., Fraiman, R. and Ransford, T. (2007) A sharp form of the Cramér-Wold theorem. *J. Theoret. Probab.*, **20**, 201–209. URL: <https://doi-org.acces.bibl.ulaval.ca/10.1007/s10959-007-0060-7>.
- Dai, B., Ding, S. and Wahba, G. (2013) Multivariate Bernoulli distribution. *Bernoulli*, **19**, 1465–1483. URL: <https://doi-org.acces.bibl.ulaval.ca/10.3150/12-BEJSP10>.
- Devroye, L., Györfi, L. and Lugosi, G. (1996) *A probabilistic theory of pattern recognition*, vol. 31 of *Applications of Mathematics (New York)*. Springer-Verlag, New York. URL: <https://doi-org.acces.bibl.ulaval.ca/10.1007/978-1-4612-0711-5>.
- Dimitrova, D. S., Kaishev, V. K. and Tan, S. (2020) Computing the Kolmogorov–Smirnov distribution when the underlying CDF is purely discrete, mixed, or continuous. *Journal of Statistical Software*, **95**, 1–42.
- Dowd, C. (2020) A new ECDF two-sample test statistic. Preprint, arXiv:2007.01360.
- Ehm, W. (1991) Binomial approximation to the Poisson binomial distribution. *Statist. Probab. Lett.*, **11**, 7–16. URL: <https://doi-org.acces.bibl>.

- ulaval.ca/10.1016/0167-7152(91)90170-V.
- Euán, C. and Sun, Y. (2020) Bernoulli vector autoregressive model. *J. Multivariate Anal.*, **177**, 104599, 19. URL: <https://doi-org.acces.bibl.ulaval.ca/10.1016/j.jmva.2020.104599>.
- Fontana, R. and Semeraro, P. (2018) Representation of multivariate Bernoulli distributions with a given set of specified moments. *J. Multivariate Anal.*, **168**, 290–303. URL: <https://doi-org.acces.bibl.ulaval.ca/10.1016/j.jmva.2018.08.003>.
- (2021) Exchangeable Bernoulli distributions: high dimensional simulation, estimate and testing. Preprint, arXiv:2101.07693.
- Fraiman, R., Moreno, L. and Ransford, T. (2022) A Cramér–Wold theorem for elliptical distributions. Preprint, arxiv.org/abs/2206.13612.
- Gardner, R. J., Gritzmann, P. and Prangenberg, D. (1996) Reconstruction of binary images from their discrete Radon transforms. In *Vision Geometry V* (eds. R. A. Melter, A. Y. Wu and L. J. Latecki), vol. 2826, 121–132. SPIE.
- Gibbs, A. L. and Su, F. E. (2002) On choosing and bounding probability metrics. *Int. Stat. Rev.*, **70**, 419–435.
- Gilbert, W. M. (1955) Projections of probability distributions. *Acta Math. Acad. Sci. Hungar.*, **6**, 195–198. URL: <https://doi-org.acces.bibl.ulaval.ca/10.1007/BF02021275>.
- Grisoni, F., Consonni, V. and Ballabio, D. (2019) Machine learning consensus to predict the binding to the androgen receptor within the CoMPARA project. *J. Chem. Inf. Model.*, **59**, 1839–1848.
- Hamedani, G. G. and Tata, M. N. (1975) On the determination of the bivariate normal distribution from distributions of linear combinations of the variables. *Amer. Math. Monthly*, **82**, 913–915. URL: <https://doi-org.acces.bibl.ulaval.ca/10.2307/2318494>.
- Heppes, A. (1956) On the determination of probability distributions of more dimensions by their projections. *Acta Math. Acad. Sci. Hungar.*, **7**, 403–410. URL: <https://doi-org.acces.bibl.ulaval.ca/10.1007/BF02020535>.
- Hoeffding, W. (1963) Probability inequalities for sums of bounded random variables. *J. Amer. Statist. Assoc.*, **58**, 13–30. URL: [http://links.jstor.org.acces.bibl.ulaval.ca/sici?sici=0162-1459\(196303\)58:301<13:PIFSOB>2.0.CO;2-D&origin=MSN](http://links.jstor.org.acces.bibl.ulaval.ca/sici?sici=0162-1459(196303)58:301<13:PIFSOB>2.0.CO;2-D&origin=MSN).
- Huber, M. and Marić, N. (2019) Admissible Bernoulli correlations. *Journal of Statistical Distributions and Applications*, **6**, 1–8.
- Irpino, A., Verde, R. and Lechevallier, Y. (2006) Dynamic clustering of histograms using Wasserstein metric. In *COMPSTAT*, 869–876. Citeseer.
- Jiang, W., Song, S., Hou, L. and Zhao, H. (2021) A set of efficient methods to generate high-dimensional binary data with specified correlation structures. *Amer. Statist.*, **75**, 310–322. URL: <https://doi-org.acces.bibl.ulaval.ca/10.1080/00031305.2020.1816213>.

- Johnson, N. L., Kotz, S. and Balakrishnan, N. (1997) *Discrete multivariate distributions*. Wiley Series in Probability and Statistics: Applied Probability and Statistics. John Wiley & Sons, Inc., New York. A Wiley-Interscience Publication.
- Kurgan, L. A., Cios, K. J., Tadeusiewicz, R., Ogiela, M. and Goodenday, L. S. (2001) Knowledge discovery approach to automated cardiac SPECT diagnosis. *Artif. Intell. Med.*, **23**, 149–169.
- López-Fidalgo, J. and Wiens, D. P. (2022) Robust active learning with binary responses. *J. Statist. Plann. Inference*, **220**, 1–14. URL: <https://doi-org.acces.bibl.ulaval.ca/10.1016/j.jspi.2022.01.004>.
- Mallows, C. L. (1972) A note on asymptotic joint normality. *Ann. Math. Statist.*, **43**, 508–515. URL: <https://doi-org.acces.bibl.ulaval.ca/10.1214/aoms/1177692631>.
- Marchetti, G. M. and Wermuth, N. (2016) Palindromic Bernoulli distributions. *Electron. J. Stat.*, **10**, 2435–2460. URL: <https://doi-org.acces.bibl.ulaval.ca/10.1214/16-EJS1175>.
- Molitor, D. and Needell, D. (2021) An iterative method for classification of binary data. *Inf. Inference*, **10**, 261–283. URL: <https://doi-org.acces.bibl.ulaval.ca/10.1093/imaiai/iaaa003>.
- Molitor, D., Needell, D., Nelson, A., Saab, R. and Salanevich, P. (2019) Classification scheme for binary data with extensions. In *Compressed sensing and its applications*, Appl. Numer. Harmon. Anal., 129–151. Birkhäuser/Springer, Cham. URL: https://doi-org.acces.bibl.ulaval.ca/10.1007/978-3-319-73074-5_4.
- Needell, D., Saab, R. and Woolf, T. (2018) Simple classification using binary data. *J. Mach. Learn. Res.*, **19**, Paper No. 61, 30.
- Oman, S. D. (2009) Easily simulated multivariate binary distributions with given positive and negative correlations. *Comput. Statist. Data Anal.*, **53**, 999–1005. URL: <https://doi-org.acces.bibl.ulaval.ca/10.1016/j.csda.2008.11.017>.
- Park, C. G., Park, T. and Shin, D. W. (1996) A simple method for generating correlated binary variates. *Amer. Statist.*, **50**, 306–310. URL: <https://doi-org.acces.bibl.ulaval.ca/10.2307/2684925>.
- Piir, G., Sild, S. and Maran, U. (2021) Binary and multi-class classification for androgen receptor agonists, antagonists and binders. *Chemosphere*, **262**, 128313.
- Rényi, A. (1952) On projections of probability distributions. *Acta Math. Acad. Sci. Hungar.*, **3**, 131–142. URL: <https://doi-org.acces.bibl.ulaval.ca/10.1007/bf02022515>.
- Teugels, J. L. (1990) Some representations of the multivariate Bernoulli and binomial distributions. *J. Multivariate Anal.*, **32**, 256–268. URL: [https://doi-org.acces.bibl.ulaval.ca/10.1016/0047-259X\(90\)90084-U](https://doi-org.acces.bibl.ulaval.ca/10.1016/0047-259X(90)90084-U).
- Xu, W., Liu, Q., Wang, Y. and Bian, X. (2021) Performance limits of one-bit compressive classification. *Signal Processing*, **178**, 107808. URL: <https://www.sciencedirect.com/science/article/pii/S0165168420303522>.

CENTRO DE MATEMÁTICA, FACULTAD DE CIENCIAS, UNIVERSIDAD DE LA REPÚBLICA,
MONTEVIDEO, URUGUAY.

Email address: `rfraiman@cmat.edu.uy`

INSTITUTO DE ESTADÍSTICA, DEPARTAMENTO DE MÉTODOS CUANTITATIVOS, FCEA,
UNIVERSIDAD DE LA REPÚBLICA, MONTEVIDEO, URUGUAY.

Email address: `leonardo.moreno@fcea.edu.uy`

DÉPARTEMENT DE MATHÉMATIQUES ET DE STATISTIQUE, UNIVERSITÉ LAVAL, QUÉBEC
CITY, CANADA.

Email address: `ransford@mat.ulaval.ca`