



Learning to increase the power of conditional randomization tests

Shalev Shaer¹ · Yaniv Romano^{1,2}

Received: 21 June 2022 / Revised: 15 November 2022 / Accepted: 11 January 2023 /
Published online: 2 February 2023

© The Author(s), under exclusive licence to Springer Science+Business Media LLC, part of Springer Nature 2023

Abstract

The model-X conditional randomization test is a generic framework for conditional independence testing, unlocking new possibilities to discover features that are conditionally associated with a response of interest while controlling type I error rates. An appealing advantage of this test is that it can work with any machine learning model to design powerful test statistic. In turn, the common practice in the model-X literature is to form a test statistic using machine learning models, trained to maximize predictive accuracy with the hope to attain a test with good power. However, the ideal goal here is to drive the model (during training) to maximize the power of the test, not merely the predictive accuracy. In this paper, we bridge this gap by introducing novel model-fitting schemes that are designed to explicitly improve the power of model-X tests. This is done by introducing a new cost function that aims at maximizing the test statistic used to measure violations of conditional independence. Using synthetic and real data sets, we demonstrate that the combination of our proposed loss function with various base predictive models (lasso, elastic net, and deep neural networks) consistently increases the number of correct discoveries obtained, while maintaining type I error rates under control.

Keywords Conditional independence testing · Conditional randomization test · Model-X Knockoffs · False discovery rate · Controlled feature selection

Editor: Manfred Jaeger.

✉ Shalev Shaer
shalev.shaer@campus.technion.ac.il

Yaniv Romano
yromano@technion.ac.il

¹ Electrical and Computer Engineering Department, Technion, Haifa, Israel

² Computer Science Department, Technion, Haifa, Israel

1 Introduction

An important task in modern data analysis is identifying a subset of explanatory features from a larger pool that are associated with a response of interest. For instance, suppose a nutritionist is interested in identifying which nutrient intakes are associated with body mass index. This information, in turn, could be leveraged to develop diets that improve human health (Chen & Li, 2013). Another example is a geneticist who is interested in finding which genetic mutations of a virus have developed resistance to a specific antiviral drug. To quote Rhee et al. (2006): “[This understanding] is essential for designing new antiviral drugs and for using genotypic drug resistance testing to select optimal therapy.” In such applications, there are often a large number of explanatory features with complex dependencies and a limited number of observations, which make the selection problem a particularly difficult task (Shah & Peters, 2020). To stress this point, consider the genetic example from above and notice that a feature selection procedure can make erroneous selections of two kinds: (i) false positives—it may select mutations that do not have any effect on drug resistance; and (ii) false negatives—it may fail to discover genetic mutations that in fact developed resistance to the antiviral drug. This paper presents a novel cost function and learning scheme whose goal is to increase the number of discoveries reported by statistical methods for feature selection, while maintaining the rate of false positives under control.

To formalize the problem, we denote by $Y \in \mathbb{R}$ the response variable and by $X = (X_1, X_2, \dots, X_d) \in \mathbb{R}^d$ the feature vector. For example, imagine that Y is a measure of drug resistance, and $X_j \in \mathbb{R}$ —the j th entry in the vector X —indicates the presence or absence of a genetic mutation in a specific location. Conditional independence testing (Candès et al., 2018; Tansey et al., 2021; Liu et al., 2020; Zhang et al., 2012) deals with the question of whether a specific feature X_j is independent of the response variable $Y \in \mathbb{R}$ after accounting for the effect of all the other features $X_{-j} = (X_1, \dots, X_{j-1}, X_{j+1}, \dots, X_d) \in \mathbb{R}^{d-1}$. That is, the null hypothesis we seek to test is

$$H_{0,j} : X_j \perp\!\!\!\perp Y \mid X_{-j}$$

against the alternative $H_{1,j} : X_j \not\perp\!\!\!\perp Y \mid X_{-j}$. In plain words, $H_{0,j}$ inquires whether the knowledge of X_j adds additional information about the response Y beyond what is already contained in X_{-j} . Consequently, we say that the j th feature is null (i.e., unimportant) if $H_{0,j}$ is true, and non-null (i.e., important) if $H_{0,j}$ is false.

Imagine we are given n data points $\{(X^i, Y^i)\}_{i=1}^n$, where $X^i \in \mathbb{R}^d$ and $Y^i \in \mathbb{R}$ are drawn i.i.d. from a joint distribution P_{XY} . Statistical tests for conditional independence leverage the observed data to formulate a decision rule on whether to reject $H_{0,j}$ and report that the feature X_j is likely to be important, or accept $H_{0,j}$ in cases where there is not enough evidence to reject the null. A *valid* test for conditional independence should control the probability of making type I error (i.e., rejecting $H_{0,j}$ when it is in fact true), by returning a p-value, a random variable that is super-uniform under the null. This p-value is then used to rigorously control the type I error rate at any user-specified level α , e.g., of 5%. The power of the test (higher is better) is the true positive rate, being the probability of correctly reporting that the non-null X_j is indeed important.

Thus far we formulated the problem of testing for a specific feature, however, in many scientific applications (e.g., genome wide association studies Sesia et al., 2020b; Benner et al., 2016) we are interested in selecting a subset of features that are associated with the response. In other words, we wish to test for all $H_{0,j}$, $j = 1, \dots, d$ *simultaneously* while controlling some statistical notion of type I error. To this end, imagine we have at hand d

p-values, one for each $H_{0,j}$, which we denote by p_j , $j = 1, \dots, d$. Recall that each p-value allows us to control the type I error for a specific hypothesis at a desired rate α . However, when testing for all the hypotheses simultaneously we must account for the multiplicity of the tests, as otherwise the probability that some of the true null hypotheses are rejected by chance alone may be excessively high; see the survey by Bender and Lange (2001) for details. Hence, we follow the seminal work by Benjamini and Hochberg (1995) and define our task as follows: identify the largest possible subset $\hat{\mathcal{S}} \subseteq \{1, \dots, d\}$ of non-null features¹ while ensuring that the false discovery rate (FDR), defined as

$$\text{FDR} = \mathbb{E} \left[\frac{|\hat{\mathcal{S}} \cap \mathcal{H}_0|}{\max\{|\hat{\mathcal{S}}|, 1\}} \right],$$

falls below a nominal level of choice $q \in (0, 1)$, e.g., $q = 0.1$ or $q = 0.2$. Above, $\mathcal{H}_0 \subseteq \{1, \dots, d\} \setminus \mathcal{H}_1$ contains the indices of the null features for which $H_{0,j}$ is true, \mathcal{H}_1 is the set of non-null features, and $|\cdot|$ returns the set-size. In words, the FDR is the expected proportion of true nulls among the rejected hypotheses. In our genetic example, the ability to control the FDR allows us to form a list of discoveries such that the majority of the selected mutations are expected to be conditionally associated with drug resistance, on average. This information may be valuable when allocating costly resources, required to further study the implications of the reported discoveries (Hawinkel et al., 2019; Manolio et al., 2009). The power of the selection procedure (larger is better) is defined as $\mathbb{E}[|\hat{\mathcal{S}} \cap \mathcal{H}_1|/|\mathcal{H}_1|]$, being the expected proportion of non-nulls that are correctly selected among all the true non-nulls. Controlling the FDR can be achieved by plugging the list of p-values p_j , $j = 1, \dots, d$ into the Benjamini-Hochberg procedure (BH) (Benjamini & Hochberg, 1995), as described in Sect. 2.2. Importantly, the power of this feature selection procedure is affected by the ability of the underlying conditional independence test to generate small p-values for the non-null features, indicating strong evidence against the corresponding null hypotheses. Our goal in this work is to increase the power of the above controlled feature selection pipeline by forming data-driven test statistics that powerfully separate non-nulls from nulls.

The model-X conditional randomization test (Candès et al., 2018) and its computationally efficient version—the holdout randomization test (HRT) (Tansey et al., 2021) that we study in this paper—are methods for conditional independence testing that gain a lot of attention in recent years. These tools are very attractive since they can leverage any powerful predictive model to rigorously test for $H_{0,j}$, under the assumption that the distribution of the explanatory features P_X is known. As such, model-X tests are especially beneficial in situations where copious amount of *unlabeled* data is available compared to labeled data, e.g., as happen in genome wide association studies (Sesia et al., 2020a, b; Bates et al., 2020). At a high level, the HRT is carried out by first splitting the data into a training set and a test set, and fit an arbitrary predictive model on the training set. Next, a dummy copy \tilde{X}_j of the original feature X_j is sampled for each sample in the test set, such that \tilde{X}_j is independent of Y given X_{-j} by construction. In our running example, one can think of \tilde{X}_j as a fake genetic mutation (sampled by a computer program) that does not carry any new information about Y if we already know all the other mutations. The test proceeds by comparing, for each observation in the test set, the accuracy of the trained predictive model applied to (i) the original feature vector X , and (ii) its modified version for which only X_j is

¹ We assume there exists a unique Markov blanket \mathcal{S} , i.e., a unique smallest subset \mathcal{S} such that Y is independent of $\{X_j\}_{j \notin \mathcal{S}}$ given $\{X_j\}_{j \in \mathcal{S}}$. We refer the reader to Candès et al. (2018, Section 2) for a detailed discussion in this regard.

replaced by \tilde{X}_j . Here, a decrease in the overall accuracy can serve as an evidence that X_j is associated with Y after accounting for the effect of X_{-j} .

In practice, the HRT is deployed by *treating the machine learning model as a black-box*, fitted to maximize predictive accuracy with the hope to attain good power. However, the ultimate objective here is to maximize the power of the test and not merely the model's accuracy, and, currently, there is no method to balance between the two. To stress this point further, there is a growing evidence that modern machine learning algorithms tend to excessively rely on spurious features—null features that are solely correlated with the non-null ones—to make accurate predictions (Sagawa et al., 2019; McCoy et al., 2019; Arjovsky et al., 2019; Peters et al., 2015; Khani & Liang, 2021). Loosely speaking, in such cases, the model tends to up-weight the importance of null features at the cost of down weighting the importance of the non-nulls. The consequence of this undesired phenomenon is that the predictive model may only show a mild drop in accuracy when replacing an important feature with its dummy copy. This renders the test to lose power, as the test accounts for the effect of the spurious correlations by design.

1.1 Our contribution

We propose novel learning schemes to fit models that are designed to *explicitly* increase the power of the HRT. This is done by ‘looking ahead’ (during training) to what will happen when the predictive model is integrated within the HRT procedure. This stands in striking contrast with the common practice, in which the model is merely trained to maximize accuracy to the extent possible. Concretely, the core of our contribution is the formulation of a novel loss function that encourages the model's accuracy to drop when replacing the original feature X_j with the dummy one \tilde{X}_j , seeking *maximum risk discrepancy* (MRD). We introduce in Sect. 3 a general, theoretically motivated, stochastic optimization procedure that enables us to augment our MRD loss to existing cost functions used in the variable selection literature, e.g., sparsity promoting penalty for deep neural networks. That section also includes a specialized learning scheme to fit sparse linear regression models (such as lasso) with our proposed loss. Extensive experiments on synthetic and real data sets are presented in Sect. 5, showing that the combination of the proposed MRD approach with existing regression algorithms *consistently improves* the power of the underlying controlled feature selection procedure. In particular, in cases where the signal is weak (i.e., in low power regimes), our method may even achieve more than 90% relative improvement compared to existing methods. Importantly, the MRD scheme preserves the validity guarantee of the p-values generated by HRT. Furthermore, we demonstrate that our new proposal is fairly robust to situations where the distribution of P_X —required to sample the dummy features—is estimated from data. These situations are of course ubiquitous in real-world applications. Lastly, a software package that implements our methods, with code for reproducing experiments, is attached to the Supplementary Material.

2 Background

2.1 Model-X randomization test

The model-X randomization test provides a valid p-value for $H_{0,j}$ in finite samples without making modeling assumptions on the conditional distribution of $Y | X$. However, this

approach assumes that the conditional distribution of $X_j | X_{-j}$ is known. The original conditional randomization test (CRT), developed by Candès et al. (2018), requires fitting a regression model many times in order to be carried out precisely. The HRT, which is extensively used in this paper, bypasses this computational issue at the expense of data splitting. Specifically, this test starts by dividing the data into disjoint training and testing sets. The training set is used to fit an arbitrary regression model \hat{f} , formulating a test statistic function $T(X_j, X_{-j}, Y; \hat{f}) \in \mathbb{R}$, e.g., the model's prediction error. Let $\{(X_i, Y_i)\}_{i \in \mathcal{I}}$ be the samples of the testing set. The HRT proceeds by repeating the following two steps for each $k = 1, \dots, K$ (treating \hat{f} as a fixed function):

- Sample a dummy feature $\tilde{X}_j^i \sim P_{X_j | X_{-j}}(X_j^i | X_{-j}^i)$, for all samples in the test set $i \in \mathcal{I}$.
- Compute the test statistic using the dummy variables

$$\tilde{t}_j^{(k)} \leftarrow \frac{1}{|\mathcal{I}|} \sum_{i \in \mathcal{I}} T(\tilde{X}_j^i, X_{-j}^i, Y^i; \hat{f}).$$

Finally, a p-value \hat{p}_j for $H_{0,j}$ is constructed by computing the empirical quantile level of the true test statistic $t^* \leftarrow \frac{1}{|\mathcal{I}|} \sum_{i \in \mathcal{I}} T(X_j^i, X_{-j}^i, Y^i; \hat{f})$ among the dummy statistics $\tilde{t}_j^{(1)}, \dots, \tilde{t}_j^{(K)}$:

$$\hat{p}_j = \frac{1 + |\{k : t^* \geq \tilde{t}_j^{(k)}\}|}{K + 1}. \quad (1)$$

Importantly, the p-value \hat{p}_j in (1) is valid, i.e., $\mathbb{P}[\hat{p}_j \leq \alpha \mid H_{0,j} \text{ is true}] \leq \alpha$. This is because (i) $\tilde{X}_j \perp\!\!\!\perp Y \mid X_{-j}$ by construction, and (ii) under the null, the random variables $t^*, \tilde{t}_j^{(1)}, \dots, \tilde{t}_j^{(K)}$ are exchangeable (Candès et al., 2018).

The validity of the p-value holds true for any choice of predictive model \hat{f} , and for any choice of test statistic $T(\cdot)$ used to compare the distributions of (Y, X_j, X_{-j}) and (Y, \tilde{X}_j, X_{-j}) . Katsevich and Ramdas (2020) proved that the likelihood ratio is the uniformly most powerful statistic, however, it requires the knowledge of the true conditional distribution of $Y \mid X$, which is unknown. In practice, a common choice for a test statistic—which we use in our proposal as well—is the squared error (Tansey et al., 2021; Bates et al., 2020) for which

$$T(X_j, X_{-j}, Y; \hat{f}) = (Y - \hat{f}(X_j, X_{-j}))^2, \quad (2)$$

where we use the notation $\hat{f}(X_j, X_{-j}) = \hat{f}(X)$ to stress that \hat{f} is applied on the vector X with its original j th feature X_j , whereas $\hat{f}(\tilde{X}_j, X_{-j})$ is applied to the same feature vector except that the original X_j is replaced by its dummy copy \tilde{X}_j . To handle the controlled feature selection problem, the HRT p-values $p_j, j = 1, \dots, d$ are most often plugged into the Benjamini-Hochberg procedure (BH) (Benjamini & Hochberg, 1995) procedure, as discussed next.

2.2 FDR control via the BH procedure

Armed with a list of d p-values p_j , one for each $H_{0,j}$, we can apply the BH procedure (Benjamini & Hochberg, 1995) to control the FDR. This method operates as follows. First, sort the given d p-values in a non-decreasing order; we denote by $p_{(k)}$ and $H_{0,(k)}$ the k th smallest p-value and its corresponding null-hypothesis, respectively. Then, for a given FDR target level q , find the largest k_0 such that $p_{(k_0)} \leq \frac{k_0}{d}q$. Lastly, reject the null hypotheses $H_{0,(1)}, H_{0,(2)}, \dots, H_{0,(k_0)}$. Importantly, the procedure described above is guaranteed to control

the FDR under the assumption that the p-values are independent. While the p-values we construct in this work do not necessarily satisfy the independence assumption, it is known that the BH procedure often controls the FDR empirically even when the p-values are dependent (except adversarial cases). This empirical observation is also endorsed in our experiments, standing in line with many other methods that plug the dependent p-values generated by CRT/HRT to BH (Tansey et al., 2021; Candès et al., 2018). With that said, it is worth noting that there exists a variant of the BH that provably controls the FDR for any arbitrary dependency structure of p_j , $j = 1, \dots, d$, however, this comes at the cost of reduced power (Benjamini & Yekutieli, 2001).

3 The proposed method

In this section, we introduce our MRD approach to fit predictive models with the goal to explicitly improve the power of the model-X tests. We begin by describing the set of ideas that construct the theoretical motivation. Then, we present a general learning scheme to fit an arbitrary predictive model (e.g., neural network) using the MRD loss (4). Lastly, we design a specialized learning scheme for sparse linear regression models (e.g., lasso).

3.1 Main idea

A central definition in our proposal is that of the risk discrepancy, expressed as

$$RD_j = \mathbb{E}[T(Y, \tilde{X}_j, X_{-j}; \hat{f}) - T(Y, X_j, X_{-j}; \hat{f})], \quad (3)$$

where \hat{f} is a fixed (pre-trained) predictive model, and $T(\cdot)$ is the squared error (2). In the context of the HRT (Sect. 2.1), a test with good power would have large and positive empirical risk discrepancy for the non-null features $j \in \mathcal{H}_1$. Therefore, during training, we wish to encourage the model \hat{f} to maximize RD_j *only for that group of features*. While this is infeasible because we do not have a priori knowledge of which of the features are in \mathcal{H}_1 , we do have access to the sampling distribution of the conditional null through the sampling of the dummy copies \tilde{X}_j . This property is formally presented in the following well-known result in the model-X literature (Candès et al., 2018), which we will use later to formulate our loss.

Proposition 1 (Candès et al., 2018) *Take $(X, Y) \sim P_{XY}$, and let \tilde{X}_{jd} be drawn independently from $P_{X_j|X_{-j}}$ without looking at Y . If $Y \perp\!\!\!\perp X_j \mid X_{-j}$, then $(Y, \tilde{X}_j, X_{-j}) \stackrel{d}{=} (Y, X_j, X_{-j})$.*

Above, the notation $\stackrel{d}{=}$ stands for equality in distribution. For completeness, the proof of the above statement is given in Supplementary Section A. A consequence of Proposition 1, important to our purposes, is given in the next corollary; see also König et al. (2021).

Corollary 1 *Let \hat{f} be a fixed predictive model, and $T(\cdot; \hat{f})$ be any test statistic. Under the setting of Proposition 1, if $(Y, \tilde{X}_j, X_{-j}) \stackrel{d}{=} (Y, X_j, X_{-j})$, then $\mathbb{E}[T(Y, \tilde{X}_j, X_{-j}; \hat{f}) - T(Y, X_j, X_{-j}; \hat{f})] = 0$.*

This motivates our formulation of a training procedure that drives the model to a solution in which the empirical risk discrepancy (evaluated on the training data) is maximized for all

features, knowing that this quantity is guaranteed to be small for the null features at test time. In fact, in Proposition 2, which is presented formally in Sect. 3.3, we prove that for a linear model with uncorrelated features, the solution that minimizes the population version of our objective function (described next) satisfies that the estimated regression coefficients of the null variables are guaranteed to be zero.

3.2 A general learning scheme

Algorithm 1 Fitting an MRD model

Input: Training data $\{(X^i, Y^i)\}_{i=1}^m$; predictive model f_θ with parameters θ ; step size η ; baseline objective hyperparameters; MRD hyperparameter λ ; number of features to generate $1 \leq N \leq d$.

- 1: **for** $k = 1, \dots, K$ **do**
- 2: Sample a random subset of features $\mathcal{P} \subseteq \{1 \dots d\}$ of the size N .
- 3: Generate a dummy \tilde{X}_j^i for each training point X_i and for each $j \in \mathcal{P}$.
- 4: Update the model parameters θ using stochastic gradient descent:

$$\theta^{k+1} \leftarrow \theta^k - \eta \nabla_\theta \left[(1 - \lambda) \mathcal{J}_{\text{base}}(\mathbf{Y}, \mathbf{X}; f_\theta) + \frac{\lambda}{|\mathcal{P}|} \sum_{j \in \mathcal{P}} \mathcal{D}(z(f_\theta), \tilde{z}(f_\theta)) \right],$$

where $z(f_\theta), \tilde{z}(f_\theta)$ stress that z, \tilde{z} are functions of θ .

- 5: **end for**

Output: A fitted model \hat{f} .

We now formalize the idea presented above. Given a training set of m i.i.d samples $\{(X^i, Y^i)\}_{i=1}^m$, our procedure begins with a generation of a dummy copy \tilde{X}_j^i for each sample and feature:

$$\tilde{X}_j^i \sim P_{X_j|X_{-j}}(X_j^i | X_{-j}^i), \quad i = 1, \dots, m, \quad j = 1 \dots d.$$

Denote the collection of training features by $\mathbf{X} \in \mathbb{R}^{m \times d}$, where X^i is the i th row in that matrix, and by $\mathbf{Y} \in \mathbb{R}^m$ a vector that contains the response variables Y^i as its entries. Let $\mathcal{J}_{\text{base}}(\mathbf{Y}, \mathbf{X}; f)$ be the objective function of the base predictive model, expressed as

$$\mathcal{J}_{\text{base}}(\mathbf{Y}, \mathbf{X}; f) = \frac{1}{m} \sum_{i=1}^m \ell(f(X^i), Y^i) + \alpha \mathcal{R}(f). \tag{4}$$

Above, the loss function ℓ measures the prediction error (MSE in our experiments), \mathcal{R} is a regularization term, and the hyperparameter α trades off the importance of the two. With this notation in place, our proposal aims to minimize the following cost function:

$$\hat{f}(x) = \underset{f \in \mathcal{F}}{\text{argmin}} (1 - \lambda) \mathcal{J}_{\text{base}}(\mathbf{Y}, \mathbf{X}; f) + \frac{\lambda}{d} \sum_{j=1}^d \mathcal{D}(z, \tilde{z}_j), \tag{5}$$

where $z = \frac{1}{m} \sum_{i=1}^m T(Y^i, X_j^i, X_{-j}^i; f)$ and $\tilde{z}_j = \frac{1}{m} \sum_{i=1}^m T(Y^i, \tilde{X}_j^i, X_{-j}^i; f)$ are the empirical risks—MSE in our context since $T(\cdot)$ is the squared error (2)—of the model evaluated on the original and synthetic triplets, respectively. The function $\mathcal{D}(z, \tilde{z}_j) \in \mathbb{R}$ measures the difference between z and \tilde{z}_j , such that $\mathcal{D}(z, \tilde{z}_j)$ gets smaller as \tilde{z}_j becomes larger than z . Therefore, (5) seeks a model \hat{f} that aligns with the base objective $\mathcal{J}_{\text{base}}$ while creating a better separation between z and \tilde{z}_j . The parameter $\lambda \in [0, 1]$ sets the relative weight between the two functionals; we present a simple, automatic approach for tuning this parameter in Supplementary Section D. In practice, we define \mathcal{D} as follows: $\mathcal{D}(z, \tilde{z}_j) = \sigma(z - \tilde{z}_j)$, where $\sigma(\cdot)$ is the Sigmoid function, i.e., $\sigma(s) = 1/(1 + e^{-s})$. Of course, other forms of this function, such as the ratio z/\tilde{z}_j , can be considered.

For ease of reference, Algorithm 1 presents a general approach for minimizing (5). In short, we use stochastic optimization and back-propagation, where each gradient step updates the model parameters by resampling \tilde{X}_j for a subset of features from $\{1, \dots, d\}$. We sample new dummies at each iteration to reduce algorithmic randomness, which, in turn, increases the stability of the procedure. In Supplementary Section B we provide a weak form of convergence of this algorithm, where the analysis requires $\mathcal{D}(\cdot)$ to be bounded from below. A requirement that is satisfied by our choice to use the Sigmoid function. Algorithm 1 is used in our experiments, implemented with a neural network as the base predictive model.

Remark 1 Algorithm 1 implements the RD penalty for all features simultaneously, thereby avoiding the need to fit a different model for each feature separately. This algorithm can be modified to fit a model with an RD penalty defined for any subset of features, and in particular to a specific feature X_j , as we demonstrate in Supplementary Section G.2.

3.3 A scheme for sparse linear regression

Next, we offer an optimization procedure tailored for the choice of sparse linear regression as the base objective, since it is widely used in the feature selection literature and often results in a powerful test (Barber & Candès, 2015; Tansey et al., 2021). To this end, consider the combination of the elastic net objective (Zou & Hastie, 2005) with our MRD loss:

$$\hat{\beta} = \underset{\beta}{\operatorname{argmin}} (1 - \lambda) \left(\overbrace{\frac{1}{2m} \|\mathbf{X}\beta - \mathbf{Y}\|_2^2 + \alpha_1 \|\beta\|_1 + \frac{\alpha_2}{2} \|\beta\|_2^2}^{\mathcal{J}_{\text{base}}(\mathbf{Y}, \mathbf{X}; f)} \right) + \frac{\lambda}{d} \sum_{j=1}^d \mathcal{D}(z(\beta), \tilde{z}_j(\beta)),$$

where $\beta \in \mathbb{R}^d$ is the regression coefficient vector, and α_1, α_2 are the elastic net penalty parameters. Above, we use the notations $z(\beta)$ and $\tilde{z}_j(\beta)$ to stress that z and \tilde{z}_j are functions of β . Using variable splitting, we can separate the sparsity promoting penalty from the remaining terms, and solve the resulting optimization problem using the alternating direction method of multipliers ADMM (Boyd et al., 2011); see Supplementary Section C for more details. Following Supplementary Algorithm 2, the advantage of the above splitting strategy is that the minimization with respect to β has a closed-form solution, being a simple projection onto the elastic net ball (Parikh & Boyd, 2014; Yu, 2013). In practice, we run this iterative algorithm until the stopping criteria suggested by Boyd et al. (2011, Section 3.3) is met.

We conclude this section by showing that under the assumptions of Proposition 2, the estimated coefficients of the null features $\hat{\beta}_j, j \in \mathcal{H}_0$ obtained by minimizing the MRD objective will be equal to zero. The proof is provided in Supplementary Section A.

Proposition 2 *Suppose that the population model is linear $Y = X^T \beta + \epsilon$, where $\mathbb{E}[X_j X_k] = 0$ for $1 \leq k \neq j \leq d$, and $\mathbb{E}[\epsilon X_j] = 0$ for all $1 \leq j \leq d$. Denote by $\hat{\beta}$ a solution to the infinite-data version of the MRD:*

$$\begin{aligned} \hat{\beta} = \operatorname{argmin}_{\beta} & (1 - \lambda) \mathbb{E}[(X^T \beta - Y)^2] \\ & + \frac{\lambda}{d} \sum_{j=1}^d \sigma(\mathbb{E}[T(Y, \tilde{X}_j, X_{-j}; \beta)] - T(Y, X_j, X_{-j}; \beta)). \end{aligned}$$

Under the above assumptions, for any $0 \leq \lambda < 1$, and for all $j \in \mathcal{H}_0$ with $\beta_j = 0$, the solution $\hat{\beta}$ satisfies that $\hat{\beta}_j = 0$ for all $j \in \mathcal{H}_0$.

4 Related work

4.1 Handling unknown conditionals

Our proposed method inherits the model X assumption that the marginal distribution P_X —required to sample \tilde{X}_j —is known: this assumption is needed to guarantee that (1) is a valid p-value (Candès et al., 2018). However, in most real-world applications P_X is unknown and thus must be estimated from the data. Here, dummy features of poor quality may break the validity of model-X randomization tests (Romano et al., 2019; Tansey et al., 2021; Sudarshan et al., 2021), including our proposal. To alleviate this concern, model X tests are equipped with various methodologies to estimate P_X , including a second-moment multivariate Gaussian estimator (Candès et al., 2018), mixture density networks (Tansey et al., 2021), generative adversarial networks (Bellet & van der Schaar, 2019), moment matching networks (Romano et al., 2019), and more (Sesia et al., 2019; Sudarshan et al., 2020; Gimenez et al., 2019; Sudarshan et al., 2021). Furthermore, model-X tests are also equipped with diagnostic tools to improve the confidence in the results obtained by this method when the conditionals are unknown. These tools include goodness-of-fit tests, designed to quantify the quality of the generated dummies, as well as controlled semi-synthetic experiments with simulated response variables that allow the user to compare the obtained empirical FDR with the desired level (Romano et al., 2019, Section 5).

Importantly, empirical evidence show that the model-X tests are robust in the sense that the FDR is often controlled in practice, especially when the testing procedure is applied with well-approximated dummy features. This is also corroborated by our experiments on real and simulated data: we practically obtain FDR control even when using the simplest Gaussian approximation to the unknown P_X , highlighting the robustness to model misspecification of the HRT in general, and our proposal, in particular.

4.2 Relevant prior art

Our proposal is connected to the two-stage “learn-then-test” approach, initiated by the influential work of Gretton et al. (2012) in the context of kernel-based tests. The original idea here is to optimize the kernel parameters in order to maximize the power of the underlying test, where ample evidence shows that this approach is extremely fruitful (Gretton et al., 2012; Liu et al., 2016; Chwialkowski et al., 2016). These contributions are in line with our proposal, suggest learning techniques to improve power, however ours is very different as it is tailored to work in combination with model-X tests.

Recently, new deep learning methods for feature selection have been proposed to improve data analysis with non-linear interactions. In particular, Borisov et al. (2019), Lemhadri et al. (2021), Yamada et al. (2020) offer sparsity-promoting layers that can be combined with arbitrary neural network architectures. These techniques can be combined with our MRD framework, as demonstrated in Sect. 5, where we utilize the CancelOut regularizer developed by Borisov et al. (2019). A different but related line of work is that of instance-wise interpretability methods, aiming at identifying a subset of relevant features for a single data point (Chen et al., 2018; Yoon et al., 2019). While this line of work is outside the scope of this paper, we believe these methods may benefit from our MRD approach. For instance, by combining the MRD with the test proposed by Burns et al. (2020), which is capable of finding instance-wise features that a given predictive model considers important.

5 Experiments

In this section, we study the performance of the proposed MRD approach in a wide variety of settings. These include experiments with linear and nonlinear simulated data of varying signal strength and sample size, model misspecification experiments, and similar experiments with semi-synthetic and real data. The conclusion of this comprehensive study is that the MRD approach *consistently* improves the power of HRT, more significantly in low power regimes.

5.1 Synthetic experiments

To evaluate the performance of our methods, we implement variable selection in a fully controlled synthetic setting in which the distribution of X and $Y | X$ are known. This experimental setting allows us to study the validity and statistical efficiency of a wide variety of predictive models \hat{f} integrated with the HRT. Following Candès et al. (2018), Romano et al. (2019), Lu et al. (2018), we generate $X \in \mathbb{R}^d$ of dimension $d = 100$ from a multivariate Gaussian distribution $\mathcal{N}(0, \Sigma)$ with $\Sigma_{i,j} = \rho^{|i-j|}$ and an auto-correlation parameter ρ . The response $Y \in \mathbb{R}$ is sampled from a sparse regression model of the form

$$Y = g(X) + \epsilon, \quad (6)$$

where $\epsilon \sim \mathcal{N}(0, 1)$ is the noise component and g is a link function. Throughout the synthetic experiments we consider the following conditional models of $Y | X$:

- **M1: linear model**, with $g(X) = X^\top \beta$, where $\beta \in \mathbb{R}^d$ has 30% randomly chosen non-zero entries of a random sign and a magnitude equal to c .

- **M2: polynomial model**, we follow Lu et al. (2018) and set a nonlinear $g(X) = (X^\top \beta)^3/2$, where β is chosen as described in **M1**.
- **M3: sum of sines** is another nonlinear model we explore, with $g(X) = \sum_{j \in \mathcal{H}_1} \sin(X_j \beta_j)$. Here, X_j and β_j are the j th entries of X and β , respectively; β is defined as in **M1**.
- **M4: interaction model**, a challenging nonlinear model with $g(X) = \sum_{j=1}^{15} X_{2j} X_{2j-1}$.

Note that for data models **M1–M3**, a variable X_j is in the non-null set $j \in \mathcal{H}_1$ if and only if $\{j = 1, \dots, d : \beta_j \neq 0\}$. In the same vein, **M4** has 30 non-null variables, those who appear in the function $g(x)$.

To apply the test, we first draw m training examples from P_{XY} , which are used to fit a predictive model of choice. We then generate an additional independent set of m test points, which are used to compute a p-value for each of the hypotheses $H_{0,j}, j = 1, \dots, d$. We use lasso, elastic net, and neural network as baseline models, where the latter is deployed only in the non-linear setting. See Supplementary Section D for additional information on the training strategy and choice of hyperparameters. We apply our MRD algorithm in combination with the above methods using the same choice of hyperparameters. In addition, we set the RD penalty parameter λ to be proportional to the validation error of the base predictive model. Supplementary Section D describes how we evaluate this error for each base model. All models are fit by normalizing the features and response variables to have a zero mean and unit variance.

5.1.1 An illustrative example

We begin with a synthetic experiment that demonstrates the effect of our proposed training scheme on the power of the selection procedure. To this end, we use $m = 400$ samples for training, fix the auto-correlation parameter ρ to be 0.1, and construct $Y | X$ that follows a polynomial data generating function (**M2**) with a signal amplitude $c = 0.14$.

As a baseline for reference, we fit a lasso regression model on the training set using ADMM (Boyd et al., 2011), where the lasso penalty parameter is tuned via cross-validation. To analyze the influence of the model on the power of the HRT, we measure the extent to which the *test* mean squared error (MSE) is increased after replacing X_j by its dummy copy \tilde{X}_j ; this difference serves as an empirical estimate of RD_j (3), which we denote by $\widehat{\text{RD}}_j$. Recall that in the context of the HRT, a positive $\widehat{\text{RD}}_j$ serves as an evidence against the null, since $\tilde{t}_j > t^*$ in (1). The left panel of Fig. 1 presents selected quantiles of $\{\widehat{\text{RD}}_j, j \in \mathcal{H}_1\}$ and $\{\widehat{\text{RD}}_j, j \in \mathcal{H}_0\}$ as a function of the ADMM iterations, where the $\widehat{\text{RD}}_j$ scores are evaluated on an independent test set. To reduce randomization, the presented $\widehat{\text{RD}}_j$ are averaged over 100 independent train and test sets. As can be seen, the quantiles that represent the non-null set increase as the training progresses, demonstrating an improvement in the sensitivity of the model to the replacement of these features by their dummy copies. Notice that the lower 25th percentile reaches a final value that is relatively close to the scores that correspond to the null set $\{\widehat{\text{RD}}_j : j \in \mathcal{H}_0\}$. Another important observation is that most of the null scores are approximately equal to zero, implying that the swap of these features with their dummy copies does not lead to significant variations in the test statistics.

Next, we fit the proposed MRD lasso model on the same 100 independent training data sets. The test error obtained by our model is similar to that of the baseline approach; both result in an averaged root mean squared error (RMSE) that is equal to 0.94. Also, similarly to the baseline model, the quantiles of the null $\widehat{\text{RD}}_j$ scores are close to zero, as shown in the right panel of Fig. 1. This phenomenon exemplifies Corollary 1. Turning to




Fig. 1 Empirical risk (MSE) difference evaluated on test data as a function of the training iterations. Left: baseline lasso model. Right: the proposed MRD method. The red (lasso) and blue (proposed method) curves represent the 25th, 50th, and 75th percentiles of the $\{\widehat{RD}_j : j \in \mathcal{H}_1\}$. The green (lasso) and orange (proposed method) color-shaded curves represent the 75th percentile of $\{\widehat{RD}_j : j \in \mathcal{H}_0\}$. Each point is evaluated by averaging \widehat{RD}_j over 100 independent data sets

the non-null set, we can see that all the quantiles of $\{\widehat{RD}_j : j \in \mathcal{H}_1\}$ obtained by our MRD model increase with the training iterations, reaching final values that are higher than the baseline approach. This advantage is corroborated by comparing the power of the selection procedure, applied with a target FDR level of $q = 0.2$. The empirical power (averaged over 100 independent experiments) obtained by the MRD model is equal to 0.521, higher than the one of the baseline model that equals 0.453. In addition, each method results in an empirical FDR that falls below the nominal rate. Here, the MRD lasso is less conservative, obtaining an FDR of 0.097 whereas the base lasso model results in 0.054.

We also present the regression coefficients obtained by lasso and MRD lasso in Fig. 6 in Supplementary Section F. In essence, the regression coefficients that correspond to the non-null features obtained by MRD lasso have higher absolute values than those obtained by lasso. By contrast, the regression coefficients that correspond to the null features are closer to zero both for lasso and MRD lasso.

5.1.2 Experiments with varying signal strength

In the illustrative example presented above we focused on a fixed signal amplitude and sample size, as well as a single base predictive model (i.e., lasso). We now turn to study the effect of the MRD scheme more systematically, by varying the signal-to-noise ratio and also include additional predictive models. We set $m = 400$ as before, but now we increase the auto-correlation parameter to $\rho = 0.25$.

We begin with the polynomial model **M2** for $Y | X$, where we control the signal-to-noise ratio by varying the signal magnitude c . The empirical power ($q = 0.2$) evaluated for each predictive model is summarized in Table 1. As displayed, our MRD approach consistently improves the power of all base models, where the gain is more significant in the low power regime. Notice that the largest improvement is achieved when using lasso as a base model and the smallest gain for the choice of a neural network model (referred as NNnet). In terms of the actual power obtained by each method, we typically receive the following relation in this experiment: MRD elastic net > MRD lasso > elastic net > lasso > MRD neural network > neural network. Supplementary Table 5 presents the FDR and root mean squared error (RMSE) obtained for each machine learning model, showing that the MRD loss tends to achieve RMSE similar to that of the base models. This table also includes comparisons to kernel ridge and random forest, concluding that

Table 1 Synthetic non-linear data with varying signal strength c

c	Lasso			Elastic Net			NNet		
	Base	MRD	% imp.	Base	MRD	% imp	Base	MRD	% imp
0.13	0.155	0.243	56.8	0.199	0.269	35.2	0.142	0.156	9.9
0.14	0.343	0.435	26.8	0.389	0.454	16.7	0.272	0.291	7.0
0.15	0.540	0.606	12.2	0.583	0.621	6.5	0.441	0.455	3.2
0.16	0.706	0.740	4.8	0.728	0.753	3.4	0.566	0.606	7.1
0.17	0.818	0.841	2.8	0.832	0.848	1.9	0.673	0.692	2.8
0.18	0.885	0.898	1.5	0.893	0.900	0.8	0.741	0.768	3.6

The table presents the empirical Power (using FDR target level $q = 0.2$) evaluated by averaging over 100 independent experiments, and the relative improvement (% imp.) of the MRD models. See Supplementary Table 5 for FDR values

our MRD elastic net is the most powerful method in this experiment. Turning to the validity of the selection procedure, following Supplementary Table 5, we observe that all methods obtain FDR below the nominal 20% level, across all data sets. Here, the MRD versions of lasso and elastic net are less conservative in the sense that they obtain higher empirical FDR compared to their base models, whereas the MRD neural network model has comparable FDR to that of its base model. An in-depth analysis of the above results is presented in Supplementary Section G.1. There, we present the empirical power and FDR as a function of the FDR target level q , as well as Q-Q plots which compare the quantiles of the produced p-values to those of a uniform distribution $U[0, 1]$.

We also perform a similar analysis to the one presented above for the sum of sines model **M3** as well as for the linear model **M1**. We summarize the results in Supplementary Section G.1 for the linear setting and in Supplementary Section G.4 for the non-linear one. The overall trend in both settings is similar to the one presented above: the MRD approach improves the power of the base models, and the FDR is controlled empirically. Lastly, Supplementary Section G.2 demonstrates that the MRD approach outperforms the baseline methods when testing for a single non-null hypothesis for a wide range of signal strength values. This experiment isolates the effect of the MRD penalty on the resulting power compared to the FDR experiments that account for multiplicity by design.

5.1.3 Experiments with varying sample size

Thus far we only considered data with a fixed number of samples. In what follows, we study the performance of the proposed MRD approach as a function of the sample size n . One can view this as a different way to modify the signal-to-noise ratio in the data, as the power is expected to increase with the sample size. We focus on the polynomial model **M2** for $Y | X$ with a correlated design in which the auto-correlation parameter $\rho = 0.25$. We generated data with a fixed signal strength $c = 0.1$, so that the power will be relatively low for the smallest sample size n . Table 2 summarizes the results, where we deployed only lasso and MRD lasso to ease with the computational load. Following that table we can see that the MRD approach consistently improves the power of the baseline model, while keeping the FDR under control.

Table 2 Synthetic experiments with non-linear data (M2) with varying number of samples n

n	MRD Lasso		Lasso		% imp. of power
	Power	FDR	Power	FDR	
1000	0.189	0.069	0.097	0.018	94.5
1500	0.403	0.109	0.288	0.048	39.9
2000	0.648	0.106	0.563	0.046	15.1
2500	0.763	0.086	0.697	0.040	9.5
3000	0.823	0.091	0.791	0.057	4.0
4000	0.927	0.116	0.921	0.061	0.7

The empirical FDR (nominal level $q = 0.2$) and power are evaluated by averaging over 50 independent experiments. All standard errors are below 0.02




Fig. 2 Synthetic experiments on a sparse interaction model (M4) with varying signal amplitude. Power and FDR ($q = 0.2$) are evaluated on 50 independent trials

5.1.4 Experiments with interaction model

In the experiments presented thus far, lasso and elastic net were more powerful than neural networks, even when applied to nonlinear data. In fact, this phenomenon is in line with the experiments reported in Tansey et al. (2021, Section 4). We now conduct experiments with a sparse interaction model for the data (M4), for which linear predictive models result in zero power in contrast to neural networks. The results are summarized in Fig. 2; we do not include linear predictive models as they did not make any discovery.² Notice that the MRD version of the neural network model achieves higher power than the base one, while controlling the FDR.

5.1.5 Experiments with cross-validation HRT

A limitation of the HRT is the reliance on data splitting, which can lead to a power loss especially when the number of samples is limited. When data is scarce, it is more

² Refer to Supplementary Section G.3 for additional details on the data generation process and the training strategy.

Table 3 Synthetic experiments with linear (**M1**) and non-linear (**M2**) data with $n < d$

Linear					
n	MRD Lasso		Lasso		% imp. of power
	Power	FDR	Power	FDR	
40	0.070	0.044	0.047	0.060	48.9
60	0.185	0.103	0.154	0.096	20.1
70	0.367	0.105	0.304	0.106	20.7
80	0.582	0.101	0.566	0.096	2.8
90	0.868	0.077	0.867	0.072	0.1
Non-Linear					
n	MRD Lasso		Lasso		% imp. of power
	Power	FDR	Power	FDR	
80	0.067	0.081	0.055	0.056	21.8
100	0.125	0.109	0.106	0.099	17.9
150	0.389	0.113	0.345	0.103	12.8
200	0.648	0.080	0.616	0.073	5.2
250	0.788	0.071	0.776	0.064	1.5

The empirical FDR (nominal level $q = 0.2$) and power are evaluated by averaging over 50 independent experiments, using CV-HRT with 8 folds. All standard errors are below 0.03

sensible to deploy the cross-validation version of the HRT (CV-HRT) proposed by Tansey et al. (2021, Section 3.1). The CV-HRT leverages the whole data for training and testing, resulting in improved power compared to a single split.

In what follows, we study the effect of the MRD approach on the CV-HRT in a small sample size regime, where $n < d$. We explore both linear (**M1**) and nonlinear (**M2**) data generating functions, where in we set $c = 1.5$ and the auto-correlation parameter $\rho = 0.25$. We implemented the CV-HRT with $K = 8$ folds and use only lasso as the base predictive model to ease the computational load. The results are summarized in Table 3, showing that (i) the power is increasing with n ; (ii) the MRD lasso consistently outperforms the baseline method; and (iii) the FDR is controlled ($q = 0.2$).

In Supplementary Section G.7 we also compare our method to the model-X knockoffs (Candès et al., 2018) and to the state-of-the-art dCRT (Liu et al., 2020) framework. The model-X knockoffs is a multiple testing framework for conditional independence which provides a finite sample FDR control under an arbitrary structure of (X, Y) . The dCRT, a test for conditional independence, is a clever technique to reduce the computational cost of the original CRT while avoiding data splitting. In a nutshell, this is done by distilling the high-dimensional information that X_{-j} contains on Y into a low-dimensional representation before applying the CRT. Yet, the complexity of dCRT is often higher than that of the CV-HRT as the former requires to fit d different predictive models—one for each leave-one-covariate-out version of the data—compared to K leave-fold-out models in the CV-HRT. Section G.7 shows that once combining CV-HRT with our MRD approach we achieve competitive power to that of the dCRT and higher power than the model-X knockoffs.

5.1.6 Experiments under model misspecification

In the experiments presented thus far we sample the dummy features \tilde{X}_j from the true distribution of $X_j | X_{-j}$. In practice, however, we often have only access to an approximation of the unknown conditionals $X_j | X_{-j}$. In Supplementary Section G.5, we include experiments under such model misspecification. There, we fit our MRD models and apply the HRT by sampling approximated dummy features from a data-driven estimation of P_X , rather than from the true one. We commence by conducting an experiment in which we control the accuracy of the estimated conditionals. This experiment demonstrates that an inflation of the empirical FDR occurs only when the sampled dummies are of poor quality, as measured by the covariance goodness-of-fit diagnostic proposed in Romano et al. (2019). Next, we study three different settings of increasing difficulty. In the first, X is generated from a correlated multivariate Gaussian, and we sample the approximated dummy features from a multivariate Gaussian whose parameters are estimated from the data. In the second, X is sampled from a Gaussian Mixture Model (GMM), whereas the dummies from a fitted multivariate Gaussian as described above. In the third, we sample X from a correlated multivariate Student t -distribution. When using a naive multivariate Gaussian fit for P_X , we report a violation in FDR control since the estimated conditionals are of low quality. However, when applying more flexible density estimation method—a Gaussian Mixture Model (Reynolds, 2009)—our method is shown to control the FDR in practice, supporting the discussion from Sect. 4.1.

5.1.7 The effect of the MRD penalty parameter

In the interest of space, we refer the reader to Supplementary Section E that illustrates the power of the test as a function of the MRD parameter λ in low, medium, and high power regimes. This section shows that our automatic approach for selecting λ achieves power that is relatively close to the best one. This section also shows that, in the high power regime (≈ 0.9), inappropriate choice of λ can reduce power compared to the base model. By contrast, in the low (≈ 0.5) and medium (≈ 0.7) power regimes, the power of the MRD approach is higher than the power of the base model (i.e., $\lambda = 0$) for all presented values of $\lambda > 0$.

5.2 Real-world application

We now apply our methods to a real-world application that already appears in the knockoff literature (Barber & Candès, 2015; Romano et al., 2019). Here, the task is to reliably detect genetic mutations associated with changes in drug resistance among human immunodeficiency viruses (HIV) of type I (Rhee et al., 2006); we study the resistance to the Lopinavir protease inhibitor drug. The response variable Y represents the log-fold increase in drug resistance measured in the i th virus, and each of the features $X_j \in \{0, 1\}$ indicates the presence or absence of a specific mutation. After applying the pre-processing steps described by Romano et al. (2019), the data contains $n = 1555$ samples with $d = 150$ features. Supplementary Section H provides additional details about this data.

Importantly, in contrast to the fully controlled setup from Sect. 5.1, here P_X is unknown and thus must be estimated. (Recall that in Supplementary Section G.5 we also study the robustness of MRD to model misspecification.) Here, we approximate P_X by fitting a multivariate




Fig. 3 Semi-synthetic experiments with real HIV mutation features and a simulated non-linear response using the generating function **M2**. Empirical power and FDR with $q = 0.2$ are evaluated over 20 random train/test splits

Gaussian on all samples $\{X^i\}_{i=1}^n$. While this estimation is most likely inaccurate due to the fact that the features are binary, we follow Romano et al. (2019) and show via semi-synthetic experiments (for which $Y | X$ is known) that the selection procedure empirically controls the FDR when sampling \tilde{X}_j from the estimated distribution. Only then, we apply the selection procedure on the real data.


Experiments with semi-synthetic data

We generate a semi-synthetic data set by simulating Y as described in (6) while treating the real X as fixed, using the polynomial non-linear generating function **M2**. Next, we randomly split the data into training and testing sets of equal size, and repeat the same experimental protocol, data normalization, and training strategy used in the synthetic experiments from Sect. 5.1. (For the neural network base model, we slightly modified the number of training epochs due to overfitting; further details are in Supplementary Section D.) Figure 3 displays the FDR and power obtained by lasso and elastic net as a function of the signal amplitude; a similar figure that describes the results obtained by the base and MRD neural network models is presented in Supplementary Section H.2. Following Fig. 3, we can see that all methods control the FDR ($q = 0.2$), demonstrating the robustness of the selection procedure to model misspecification. Regarding efficiency, our proposal consistently improves the power of the base models. In contrast to the fully synthetic experiments (see Table 1), here the gap between the power of the base model and its MRD version increases as a function of the signal amplitude. Additional semi-synthetic experiments with a linear model that also explore the effect of the sample size are presented in Supplementary Sections H.3 and H.4. Both demonstrate that a similar gain in performance is consistently obtained by our MRD method.

Experiments with full real data

Next, we perform variable selection on the real data, i.e., with the real (X, Y) pairs. Here, we only use lasso and elastic net as base models, and deploy the same training strategy and data-normalization step described in the semi-synthetic experiments. We repeat the variable selection pipeline for 1000 independent train/test splits and present the results in Fig. 4. As portrayed, our MRD models lead to more findings than the baseline methods for various FDR levels, which aligns with the semi-synthetic experiment results. Here, this gap increases with the target FDR level.

Fig. 4 The average number of drug-resistance mutations in the HIV discovered by base/MRD models



6 Discussion

This paper presents a novel training framework to increase the power of the model-X randomization test. Our method has an advantage over the common practice for which the model-fitting step is not tailored to maximize the power of the HRT. Focusing on FDR control, through experiments, we demonstrate that our method consistently outperforms a wide range of base predictive models in various settings.

This paper advocates conditional independence testing as a statistical approach to avoid the selection of “spurious features”, i.e., null features that are only correlated with the non-null ones. Yet, conditional independence testing has an inherent limitation: it is fundamentally impossible to distinguish nearly identical features. Put differently, it inevitable that we would suffer from low power when analyzing data with extremely correlated features. Figure 19 illustrates this phenomenon, where we apply the HRT for multivariate Gaussian X with increasing correlation among the features, controlled by the auto-correlation coefficient ρ ; see Supplementary Section G.6 for more details. As portrayed, the power of the test is reduced with the increase of ρ , making the selection procedure fundamentally harder. Observe that MRD lasso and MRD elastic net yield higher power than their vanilla counterparts, where the gap between them is reduced as the problem becomes fundamentally harder.

One way to deal with the hardness of this problem is to group (or even prune) features that are nearly identical in a given data. For example, Sesia et al. (2019) apply hierarchical clustering to identify groups of features in such a way that the correlation between the features across the two clusters is not too high. Then, one can apply the selection procedure only on the cluster representatives, which are substantially less correlated by design. This strategy increases power, however at the cost of reducing the resolution of the findings (Sesia et al., 2019). Rather than testing at the level of individual features, the test is applied to groups of features among which there may be one or more important variables. We believe it will be of great importance to continue this line of work in combination with our MRD approach, and offer practical tools that allow the use of conditional independence testing even for highly correlated data.

In a broader view of the controlled variable selection problem, the model-X knock-off filter (Candès et al., 2018) provides a rigorous FDR control in finite samples under an arbitrary dependency structure of (X, Y) . This stands in contrast with the BH procedure that provably controls the FDR under independence of the p -values constructed by the

HRT; see discussion in Sect. 2.2. With that said, sampling the knockoff dummies is more challenging and the model fit must be done by augmenting the original features and their knockoff dummies, which increases the dimensions of the problem. Therefore, extending our proposal to work with the knockoffs framework is not straightforward (to say the least), but, at the same time, of great interest.

Appendix A Supplementary proofs

Proof of Proposition 1 For completeness, we prove Proposition 1 by following the steps from Romano et al, 2020, Proposition 1. We use discrete random variables for simplicity. Denote by $x' = x_{-j}$ and observe that

$$\begin{aligned} \mathbb{P}(Y = y, X_j = x_j, X_{-j} = x') &= \mathbb{P}(Y = y, X_j = x_j \mid X_{-j} = x') \cdot \mathbb{P}(X_{-j} = x') \\ (\text{Since } Y \perp\!\!\!\perp X_j \mid X_{-j}) &= \mathbb{P}(Y = y \mid X_{-j} = x') \cdot \mathbb{P}(X_j = x_j \mid X_{-j} = x') \\ &\quad \cdot \mathbb{P}(X_{-j} = x') \\ (\text{By construction of } \tilde{X}_j) &= \mathbb{P}(Y = y \mid X_{-j} = x') \cdot \mathbb{P}(\tilde{X}_j = x_j \mid X_{-j} = x') \\ &\quad \cdot \mathbb{P}(X_{-j} = x') \\ (\text{Since } Y \perp\!\!\!\perp \tilde{X}_j \mid X_{-j}) &= \mathbb{P}(Y = y, \tilde{X}_j = x_j \mid X_{-j} = x') \cdot \mathbb{P}(X_{-j} = x') \\ &= \mathbb{P}(Y = y, \tilde{X}_j = x_j, X_{-j} = x'). \end{aligned}$$

□

Proof of Proposition 2 Following Corollary 1, for any null feature $\beta_j = 0$, we know that the following holds

$$\mathbb{E}[T(Y, \tilde{X}_j, X_{-j}; \hat{\beta}) - T(Y, X_j, X_{-j}; \hat{\beta})] = 0,$$

for any value of $\hat{\beta}_j$. We now turn to study the effect of $\hat{\beta}_j, j \in \mathcal{H}_0$ on a non-null feature $k \in \mathcal{H}_1$, by re-writing the test statistic:

$$\begin{aligned} \mathbb{E}[T(Y, X_k, X_{-k}; \hat{\beta})] &= \mathbb{E}[(Y - X_j \hat{\beta}_j - X_k \hat{\beta}_k - X_{-(j,k)}^T \hat{\beta}_{-(j,k)})^2] \\ &= \mathbb{E}[(X_j \beta_j + X_k \beta_k + X_{-(j,k)}^T \beta_{-(j,k)} \\ &\quad + \epsilon - X_j \hat{\beta}_j - X_k \hat{\beta}_k - X_{-(j,k)}^T \hat{\beta}_{-(j,k)})^2] \\ &= \mathbb{E}\left[(X_k(\beta_k - \hat{\beta}_k) + X_{-(j,k)}^T(\beta_{-(j,k)} \right. \\ &\quad \left. - \hat{\beta}_{-(j,k)}) + \epsilon - X_j \hat{\beta}_j\right]^2 \tag{A1} \\ &= \mathbb{E}[(X_k(\beta_k - \hat{\beta}_k))^2] + \mathbb{E}[(X_{-(j,k)}^T(\beta_{-(j,k)} - \hat{\beta}_{-(j,k)}))^2] \\ &\quad + \mathbb{E}[\epsilon^2] + \hat{\beta}_j^2 \mathbb{E}[X_j^2]. \end{aligned}$$

The third equality uses the fact that $\beta_j = 0$ and the last one relies on the assumptions that $\mathbb{E}[X_j, X_k] = 0$ for all $1 \leq j \neq k \leq d$ and $\mathbb{E}[X_j, \epsilon] = 0$ for all $1 \leq j \leq d$. Similarly, since $\mathbb{E}[X_j, \tilde{X}_k] = 0$ we have also that

$$\mathbb{E}[T(Y, \tilde{X}_k, X_{-k}; \hat{\beta})] = \mathbb{E}[(X_k \beta_k - \tilde{X}_k \hat{\beta}_k)^2] + \mathbb{E}[(X_{-(j,k)}^T (\beta_{-(j,k)} - \hat{\beta}_{-(j,k)}))^2] + \mathbb{E}[\epsilon^2] + \hat{\beta}_j^2 \mathbb{E}[X_j^2].$$

As a result, we get that $\hat{\beta}_j, j \in \mathcal{H}_0$ does not influence the risk discrepancy of the null and non null features. To complete the proof, we now show that any choice of $\hat{\beta}_j \neq 0$ can only increase the MSE term in the objective function. Applying similar steps as in (A1), we get

$$\mathbb{E}[(Y - X^T \hat{\beta})^2] = \mathbb{E}[(X_{-j}^T (\beta_{-j} - \hat{\beta}_{-j}))^2] + \mathbb{E}[\epsilon^2] + \hat{\beta}_j^2 \mathbb{E}[X_j^2],$$

which is minimized for the choice of $\hat{\beta}_j = 0$. □

Appendix B Supplementary analysis of the optimization algorithm

In this section, we follow Romano et al. (2019), Sanjabi et al. (2018) and provide a weak form of convergence to the general learning scheme presented in Algorithm 1 of the main manuscript. Denote by f_{θ_k} a predictive model, parameterized by θ_k (e.g., the weights and biases of a neural network), at step k of the optimization procedure. Recall that in Algorithm 1, in each iteration we choose a random subset of indices \mathcal{P}_k and generate a *fresh* dummy copy \tilde{X}_j^i of X_j^i for all $j \in \mathcal{P}_k$. For ease of notation, we denote by $\tilde{\mathbf{X}}_k \in \mathbb{R}^{m \times |\mathcal{P}_k|}$ the dummy matrix with the new \tilde{X}_j^i as its (i, j) 's entry. Thus, conditional on $(\mathcal{P}_k, \tilde{\mathbf{X}}_k)$ the objective function

$$\mathcal{J}(\mathbf{X}, \mathbf{Y}, f_{\theta_k}, \mathcal{P}_k, \tilde{\mathbf{X}}_k) = (1 - \lambda) \mathcal{J}_{\text{base}}(\mathbf{Y}, \mathbf{X}; f_{\theta_k}) + \frac{\lambda}{|\mathcal{P}_k|} \sum_{j \in \mathcal{P}_k} \mathcal{D}(z(f_{\theta_k}), \tilde{z}_j(f_{\theta_k})),$$

is a deterministic function of θ_k . Now, we can define

$$\mathcal{J}_{\theta_k} = \mathbb{E}_{\mathcal{P}_k, \tilde{\mathbf{X}}_k} [\mathcal{J}(\mathbf{X}, \mathbf{Y}, f_{\theta_k}, \mathcal{P}_k, \tilde{\mathbf{X}}_k) \mid \theta_k] \tag{B2}$$

and denote by $\nabla \mathcal{J}_{\theta_k}$ its gradient with respect to θ_k . As depicted in Algorithm 1, we estimate $\nabla \mathcal{J}_{\theta_k}$ by sampling one realization of \mathcal{P}_k and $\tilde{\mathbf{X}}_k$:

$$g_t = \nabla \mathcal{J}(\mathbf{X}, \mathbf{Y}, f_{\theta_k}, \mathcal{P}_k, \tilde{\mathbf{X}}_k), \tag{B3}$$

resulting in an unbiased estimate of $\nabla \mathcal{J}_{\theta_k}$. Since all random variables are observed at the k th step, g_t from (B3) is deterministic and backpropogation can be applied to update $\theta_{k+1} \leftarrow \theta_k - \eta g_t$, where η is the learning rate (step size). Our analysis assumes that there exists a finite Lipschitz constant L , which satisfies the following relation for all θ', θ'' :

$$\|\nabla \mathbb{E}_{\mathcal{P}, \tilde{\mathbf{X}}} [\mathcal{J}(\mathbf{X}, \mathbf{Y}, f_{\theta}, \mathcal{P}, \tilde{\mathbf{X}}) \mid \theta'] - \nabla \mathbb{E}_{\mathcal{P}, \tilde{\mathbf{X}}} [\mathcal{J}(\mathbf{X}, \mathbf{Y}, f_{\theta}, \mathcal{P}, \tilde{\mathbf{X}}) \mid \theta'']\|_2 < L \|\theta' - \theta''\|_2.$$

Denote $\Delta = \frac{2}{L} \sup (\mathcal{J}_{\theta_1} - \mathcal{J}^{\#})$, where \mathcal{J}_{θ_1} is the objective obtained by the initialized model, i.e., by assigning θ_1 in (B2), and $\mathcal{J}^{\#}$ is a lower bound uniformly the value of \mathcal{J}_{θ_k} for all k .

The supremum is taken over all possible values of \mathcal{P} , $\tilde{\mathbf{X}}$ and θ_1 . Suppose we can uniformly bound from below the RD term, which holds for our choice of \mathcal{D} from Sect. 3.2, since $\mathcal{D}(z, \tilde{z}_j) = \sigma(z - \tilde{z}_j) > 0$. Suppose further that we can uniformly lower bound the objective of the base predictive model; for example, the MSE loss is bounded from below by 0. With these notations in place, we can immediately invoke Theorem 2 from Romano et al. (2019), which is stated below formally for convenience. In short, Theorem 3 provides us a weak form of convergence as it shows that the average squared norm of the gradient $\mathbb{E}[\|\nabla \mathcal{J}_{\theta_k}\|_2^2 \mid \theta_1]$ decreases as $\mathcal{O}(1/\sqrt{K})$; a detailed discussion on this result is provided by Romano et al. (2019); see also Sanjabi et al. (2018).

Theorem 3 (Theorem 2, Romano et al. 2019) *Consider a fixed training set \mathbf{X}, \mathbf{Y} , and use the notation from above. Assume*

$$\mathbb{E}[\|g_t - \nabla \mathcal{J}_{\theta_k}\|_2^2 \mid \theta_k] \leq \sigma^2, \quad \forall k \leq K,$$

for some $\sigma \in \mathbb{R}$. Then for any initial θ_1 and a suitable value of Δ ,

$$\frac{1}{K} \sum_{k=1}^K \mathbb{E}[\|\nabla \mathcal{J}_{\theta_k}\|_2^2 \mid \theta_1] \leq \frac{1}{K} \frac{L\Delta}{\eta(2-L\eta)} + \frac{L\sigma^2\eta}{(2-L\eta)}.$$

Specifically, choosing $\eta = \min\{\frac{1}{L}, \frac{\eta_0}{\sigma\sqrt{K}}\}$ for some $\eta_0 > 0$ gives

$$\frac{1}{K} \sum_{k=1}^K \mathbb{E}[\|\nabla \mathcal{J}_{\theta_k}\|_2^2 \mid \theta_1] \leq \frac{L^2\Delta}{K} + \left(\eta_0 + \frac{\Delta}{\eta_0}\right) \frac{L\sigma}{\sqrt{K}}.$$

Appendix C Supplementary details on fitting sparse linear regression models

Algorithm 2 presents our specialized learning scheme for linear regression model. This algorithm minimizes the following objective:

$$\hat{\beta} = \underset{\beta, \mathbf{v}}{\operatorname{argmin}} (1 - \lambda) \left(\frac{1}{2m} \|\mathbf{X}\mathbf{v} - \mathbf{Y}\|_2^2 + \alpha_1 \|\beta\|_1 + \frac{\alpha_2}{2} \|\beta\|_2^2 \right) + \frac{\lambda}{d} \sum_{j=1}^d \mathcal{D}(z(\mathbf{v}), \tilde{z}_j(\mathbf{v}))$$

subject to $\beta = \mathbf{v}$,

using ADMM (Boyd et al., 2011).

Algorithm 2 Fitting MRD elastic net via ADMM

Input: Training data $\{(X^i, Y^i)\}_{i=1}^m$; elastic net hyperparameters α_1, α_2 ; ADMM hyperparameter ρ ; MRD hyperparameter λ ; number of features to generate $1 \leq N \leq d$.

- 1: Initialize $\beta^0 \leftarrow \mathbf{0}, \mathbf{v}^0 \leftarrow \mathbf{0}, \mathbf{u}^0 \leftarrow \mathbf{0}$; where \mathbf{u} is the scaled Lagrange multipliers vector.
- 2: **for** $k = 1, \dots, K$ **do**
- 3: Sample a random subset of features $\mathcal{P} \subseteq \{1 \dots d\}$ of the size N .
- 4: Generate a dummy \tilde{X}_j^i for each training point X_i and for each $j \in \mathcal{P}$.
- 5: Update \mathbf{v} using gradient descent and back-propagation:

$$\mathbf{v}^{k+1} \leftarrow \underset{\mathbf{v}}{\operatorname{argmin}} \frac{1 - \lambda}{2m} \|\mathbf{X}\mathbf{v} - \mathbf{Y}\|_2^2 + \frac{\lambda}{|\mathcal{P}|} \sum_{j \in \mathcal{P}} \mathcal{D}(z(\mathbf{v}), \tilde{z}_j(\mathbf{v})) + \frac{\rho}{2} \|\mathbf{v} - \beta^k + \mathbf{u}^k\|_2^2.$$

- 6: Update β using the proximal mapping of elastic net:

$$\beta^{k+1} \leftarrow \frac{1}{1 + \alpha_2} S_{\alpha_1/\rho}(\mathbf{v}^{k+1} + \mathbf{u}^k/\rho),$$

where $S_\tau(x) = \max\{0, x - \tau\} - \max\{0, -x - \tau\}$.

- 7: Update \mathbf{u} by setting $\mathbf{u}^{k+1} \leftarrow \mathbf{u}^k + \mathbf{v}^{k+1} - \beta^{k+1}$.
- 8: **end for**

Output: A regression coefficient vector $\hat{\beta}$.

Appendix D Supplementary details on predictive models and training strategy

Section 5 of the main manuscript compares the performance of several regression models on synthetic and real data sets. Below, we provide additional details on each method.

- **Lasso:** we fit a lasso regression model using Python’s `sklearn` package and tuned the lasso penalty parameter via 5-fold cross-validation.
- **Elastic net:** an elastic net model is fitted using the same Python package and training strategy described above.
- **Neural network (NNet):** the network architecture consists of two fully connected layers with a hidden dimension of 16; the model is fitted using Adam optimizer (Kingma & Ba, 2014) with a fixed learning rate of 0.005. To avoid over-fitting, a dropout regularization (Srivastava et al., 2014) is deployed, with a rate of 0.5. We also use CancelOut (Borisov et al., 2019) regularization to promote sparsity in the entry layer.³ In Sect. 5.1, we train the network for 60 epochs, whereas in Sect. 5.2 we reduced this number due to over-fitting.

³ The implementation is available online at <https://github.com/unnir/CancelOut>




Fig. 5 The effect of the MRD regularization on power, demonstrated using a synthetic data with correlated multivariate Gaussian features and a simulated response that follows a non-linear model. Each panel displays the empirical power (evaluated by averaging over 50 independent experiments) as a function of the MRD penalty parameter λ , using lasso as the base predictive model. The green vertical line represents the value obtained by our proposed automatic approach for choosing λ . Left: low power regime for which the signal amplitude $c = 0.14$. Middle: low power regime for which the signal amplitude $c = 0.15$. Right: high power regime for which the signal amplitude $c = 0.18$

Specifically, we use 15 epochs for $c \in \{0.5, 0.625, 0.75\}$, and 10 epochs for $c \in \{1, 1.25\}$, where c is the signal strength of the semi-synthetic data.

- **MRD lasso:** we fit the model by minimizing (5) in combination with lasso regularization using Algorithm 2. The lasso penalty parameter is equal to the one used in **lasso**. The MRD penalty λ is set to be equal to $\min\{0.8, 0.8 \cdot \text{MSE} - \text{Validation}\}$, where $\text{MSE} - \text{Validation}$ is the average MSE evaluated on the 5 folds used to tune the hyperparameter of the baseline lasso model. Regarding the hyperparameter of the ADMM framework, we set $\rho = 1$ and the stopping criteria parameters $\epsilon^{\text{rel}} = 10^{-3}$, $\epsilon^{\text{abs}} = 5 \cdot 10^{-4}$; we deployed the stopping rule suggested in (Boyd et al, 2011, Section 3.3).
- **MRD elastic net:** here, we combine the elastic net penalty with our RD loss and use the same training strategy described in **MRD lasso**.
- **MRD neural network:** we choose the same architecture, optimizer, and hyperparameters outlined in **neural network**, and fit the model as described in Algorithm 1. We set the MRD penalty λ to be equal to $\min\{0.8, 0.8 \cdot \text{MSE} - \text{Validation}\}$, where the $\text{MSE} - \text{Validation}$ is evaluated as follows. First, we split the training data into training (80%) and validation (20%) sets. Then, we fit the base neural network model (without the penalty) on the training data, and use this model to compute the $\text{MSE} - \text{Validation}$ on the validation set.

We run all the numerical experiments on our local cluster, which consists of two AMD servers with 24 CPU cores each.

Appendix E Supplementary details on the effect of λ

Section D of the Supplementary Material presents a simple way to choose the penalty parameter λ . Here, we demonstrate the impact of different values of λ on the performance of the MRD approach. Figure 5 presents the empirical power obtained by using the MRD lasso, as a function of the penalty parameter λ . The vertical green line represents the choice of our approach for choosing λ . Here, we use the same non-linear synthetic data as in

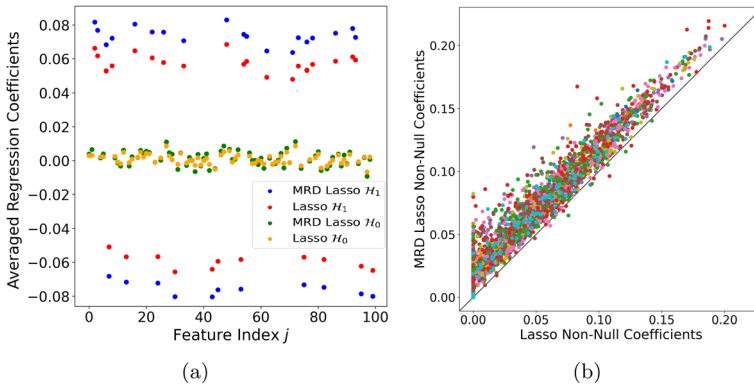


Fig. 6 Comparison between MRD lasso regression coefficients and vanilla lasso regression coefficients. The predictive models are fitted on the synthetic non linear data from Sect. 5.1.1. Left: the regression coefficients of lasso and MRD lasso for all features. Each point in this panel is obtained by averaging the estimated regression coefficient over 100 independent trials. Green and orange dots represent null features; blue and red dots represent non-null features. Right: scatter plot of the magnitude of MRD lasso regression coefficients as a function of the vanilla lasso regression coefficients for non-null features. Each color represent a different non-null feature (Color figure online)

Sect. 5.1.1, with three different signal amplitude values, $c = 0.14$, $c = 0.15$ and $c = 0.18$, representing low, medium and high power regimes. As can be seen, in the low and medium power regimes, the choice of every presented value of λ results in higher power than the base model (i.e., $\lambda = 0$). However, in the high power regime, higher values of λ reduces the power of the base models. This phenomenon corroborates our approach to choose λ to be proportional to the prediction error of the base model, as presented in Supplementary Section D.

Following that figure, we can also see that our proposed automatic approach for choosing λ , returns a penalty parameter that is fairly close the best one.

Appendix F The effect of MRD on lasso’s regression coefficients

In Sect. 3.3 we describe our learning procedure for the choice of sparse linear regression as the base objective. Here, we demonstrate the impact of the MRD approach on the estimated regression coefficients. Figure 6a presents the average of the estimated regression coefficients of lasso and MRD lasso from the experiment presented in Sect. 5.1.1. As can be seen, on average, the estimated regression coefficients of the null features that correspond to the MRD lasso are slightly larger than those of lasso, but both are relatively close to zero. By contrast, the coefficients of the non-null features have a larger magnitude, where it is evident that the MRD approach further increases the coefficients’ magnitude. Figure 6b also presents the magnitude of the estimated coefficients of MRD lasso that correspond to the non-null features as a function of lasso coefficients. Each color is for a different important feature, so there are 100 points per feature, one for a different trial. We can see that the MRD lasso’s coefficients are with larger magnitude than those of lasso, where the amplification of the lower coefficients is more significant than of the higher ones. Specifically,

it can be observed that in several cases lasso nullifies coefficients of important features, whereas MRD lasso yields non-zero values (the opposite is not true).

Appendix G Supplementary details on synthetic experiments

G.1 Further Analysis of The Numerical Experiments

Linear $Y | X$ and additional comparisons

In Table 1 of the main manuscript we compared the performance of MRD lasso, MRD elastic net and MRD NNet to their base models by varying the signal strength while fixing the target FDR level to $q = 0.2$. For completeness, we provide here the root mean squared error (RMSE) and the empirical FDR obtained by each model as well as additional comparisons to other baseline methods. We also present here the empirical power, FDR and RMSE of our experiments in the linear setting, where $Y | X$ follows **M1**. Table 4 focuses on the linear setting, where we include a ridge regression model in addition to lasso and elastic net. The penalty parameter of the ridge regression model is tuned via 5-fold cross-validation. Following that table, we can see that the empirical power obtained by the ridge model does not exceed the ones achieved by our proposed MRD methods. Observe also that the RMSE scores obtained by MRD lasso and MRD elastic net are similar to that of the corresponding baseline models. Turning to the non-linear case, Table 5 summarizes the performance of random forest and kernel ridge regression in addition to lasso, elastic net, and neural network. We fit the kernel ridge regression model using the `sklearn` software package, where we choose a polynomial kernel of degree 3 and a fixed penalty parameter that equal to 1. The random forest model is also fitted using `sklearn`, where we choose the software's default hyperparameters except the number of trees and the maximum depth of each tree; we set both to 100. In terms of power, the kernel ridge model outperforms the baseline models, however, it does not exceed the performance of MRD lasso and MRD elastic net. The random forest model is less competitive and yields lower power compared to the other methods, across all experiments. Regarding the RMSE, overall we can see that the MRD models perform similarly to their baseline counterparts, except the case of a neural network base model. There, the RMSE of the MRD neural network is slightly improved.

Power as a function of the target FDR

The experiments presented above focus on feature selection with a fixed target FDR level. Figures 7, 8 present additional comparisons in which we vary the target FDR level while keeping the signal amplitude fixed. Following Fig. 7, which corresponds to the linear case (i.e., $Y | X$ follows **M1**), we can see that the MRD models outperform their baseline counterparts in terms of power. Observe that the FDR is controlled in all cases, where the base models are more conservative than their MRD versions. A similar trend is illustrated in Fig. 8 that represents the nonlinear setting (where $Y | X$ follows **M2**). Here, we include a neural network model in addition to the sparse linear methods. Notice that the MRD neural network model attains better power than the baseline method, while the empirical FDR of the two is similar, both fall below the target level.

p-value analysis

Herein, we study the behaviour of the p-values produced by the MRD approach. Figure 12 presents Q-Q plots of the p-values from the experiments summarized in Table 1, focusing on the case where $c = 0.14$ with lasso as the baseline model. The left panel of

Table 4 Synthetic linear data with varying signal strength c

c	MRD Lasso			Lasso			MRD Elastic Net			Elastic Net			Ridge		
	Power	FDR	RMSE	Power	FDR	RMSE	Power	FDR	RMSE	Power	FDR	RMSE	Power	FDR	RMSE
0.11	0.457	0.091	1.021	0.377	0.050	0.998	0.502	0.095	0.996	0.436	0.073	0.991	0.467	0.135	0.979
0.12	0.578	0.097	0.995	0.514	0.053	0.977	0.604	0.095	0.975	0.560	0.073	0.972	0.578	0.141	0.955
0.13	0.679	0.091	0.964	0.628	0.057	0.952	0.688	0.102	0.950	0.656	0.072	0.947	0.665	0.139	0.931
0.14	0.752	0.085	0.931	0.722	0.058	0.923	0.766	0.095	0.922	0.739	0.074	0.920	0.742	0.144	0.907
0.15	0.821	0.095	0.897	0.790	0.059	0.892	0.823	0.094	0.892	0.805	0.072	0.891	0.804	0.145	0.884
0.16	0.867	0.083	0.865	0.855	0.061	0.862	0.871	0.098	0.863	0.858	0.072	0.862	0.850	0.145	0.860
0.17	0.901	0.086	0.837	0.891	0.064	0.835	0.903	0.093	0.836	0.895	0.075	0.835	0.888	0.142	0.837
0.18	0.933	0.082	0.811	0.925	0.069	0.810	0.932	0.098	0.811	0.926	0.073	0.811	0.917	0.143	0.815

The table presents the empirical Power (using FDR target level $q = 0.2$) evaluated by averaging over 100 independent experiments, and the relative improvement (% imp.) of the MRD models. Standard errors are below 0.012

Table 5 Complementary results to those presented in Table 1. All standard errors are below 0.02

c	MRD Lasso			Lasso			MRD Elastic Net			Elastic Net		
	Power	FDR	RMSE	Power	FDR	RMSE	Power	FDR	RMSE	Power	FDR	RMSE
0.13	0.243	0.095	0.979	0.155	0.04	0.973	0.269	0.098	0.964	0.199	0.048	0.967
0.14	0.435	0.093	0.949	0.343	0.039	0.946	0.454	0.1	0.938	0.389	0.062	0.941
0.15	0.606	0.09	0.916	0.54	0.049	0.914	0.621	0.105	0.908	0.583	0.063	0.91
0.16	0.74	0.092	0.882	0.706	0.053	0.879	0.753	0.105	0.878	0.728	0.074	0.879
0.17	0.841	0.093	0.851	0.818	0.057	0.849	0.848	0.098	0.849	0.832	0.075	0.849
0.18	0.898	0.093	0.823	0.885	0.06	0.821	0.9	0.101	0.822	0.893	0.075	0.822

c	MRD NNet			NNet			Kernel Ridge			Random Forest		
	Power	FDR	RMSE	Power	FDR	RMSE	Power	FDR	RMSE	Power	FDR	RMSE
0.13	0.156	0.111	0.994	0.142	0.131	1.018	0.204	0.16	0.988	0.044	0.149	0.993
0.14	0.291	0.119	0.964	0.272	0.132	0.984	0.383	0.134	0.956	0.087	0.109	0.986
0.15	0.455	0.124	0.933	0.441	0.14	0.948	0.551	0.139	0.923	0.141	0.14	0.979
0.16	0.606	0.133	0.902	0.566	0.148	0.915	0.714	0.144	0.892	0.21	0.131	0.973
0.17	0.692	0.137	0.873	0.673	0.143	0.885	0.815	0.139	0.864	0.26	0.157	0.968
0.18	0.768	0.144	0.849	0.741	0.137	0.859	0.872	0.137	0.84	0.302	0.14	0.964




Fig. 7 Synthetic experiments with correlated multivariate Gaussian features and a simulated response that follows a linear model with a constant signal strength $c = 0.12$. Each graph displays the empirical FDR and power (evaluated by averaging over 100 independent experiments) as a function of the target FDR level q . Left: MRD lasso and lasso. Right: MRD elastic net and elastic net




Fig. 8 Synthetic experiments with correlated multivariate Gaussian features and a simulated response that follows a non-linear model. The signal strength is fixed and equals to $c = 0.14$. Left: MRD lasso and lasso. Middle: MRD elastic net and elastic net. Right: MRD neural network (NNet) and NNet. The other details are as in Fig. 7




Fig. 9 Testing for a single hypothesis on synthetic data with correlated multivariate Gaussian features and a simulated response that follows a linear model. The significance level is fixed and equals to $\alpha = 0.05$. The empirical power is presented as a function of the signal strength c , and evaluated by averaging over 100 independent experiments. Left: MRD lasso and lasso. Right: MRD elastic net and elastic net

Fig. 12 presents the empirical quantiles of the p-values that correspond to the null features as a function of the quantiles of the uniform distribution on the $[0, 1]$ segment. As can be seen, the HRT p-values for the null features are valid but conservative both for lasso and MRD lasso where the distribution of the latter p-values are closer to the uniform distribution. The right panel of Fig. 12 presents an analogous plot but for the non-null features. Here, when combining HRT with MRD lasso we get p-values with lower quantiles than the ones of lasso, and both lie below the quantiles of the uniform distribution. Importantly, the MRD lasso leads to smaller p-values than lasso for the non-null features, demonstrating the advantage of the MRD approach.

G.2 Testing for a single hypothesis

We now analyze the effect of the MRD approach when testing for a single non-null hypothesis. Specifically, we choose at random a non-null j from \mathcal{H}_1 , and test for $H_{0,j}$ at level $\alpha = 0.05$. We first consider the p-value evaluated via the HRT, using the same models and data from Tables 4, 5. Figure 9 focuses on the case for which $Y | X$ follows a linear model (i.e., **M1**), presenting the power of lasso, elastic net, and their MRD versions as a function of the signal strength c . As can be seen, the results are aligned with the FDR experiments in the sense that the MRD approach increases the power of the test. Figure 10 repeats the same experiment in the non-linear case (where $Y | X$ follows **M2**). Here, once more, an improvement in power is achieved by the MRD approach. Notice that the gain in performance is not consistent for the base neural network model.

Next, we demonstrate the advantage of the MRD when optimizing for a specific feature j . To this end, we fit a model via Algorithm 1 but set $\mathcal{P} = \{j\}$ in line 2 of Algorithm 1. We use the same synthetic experimental setting with a non linear data model (**M2**) as in the experiments from Sect. 5.1.2. Specifically, we fit MRD lasso and MRD elastic net for a single arbitrary non-null feature instead of all the features simultaneously. Figure 11 presents the power of lasso, MRD lasso, elastic net, and MRD elastic net as a function of the signal strength c . As can be seen, the power of the MRD models is higher than that of the baseline models.




Fig. 10 Testing for a single hypothesis on synthetic data with correlated multivariate Gaussian features and a simulated response that follows a non-linear model. Left: MRD lasso and lasso. Middle: MRD elastic net and elastic net. Right: MRD neural network (NNet) and NNet. The other details are as in Fig. 9

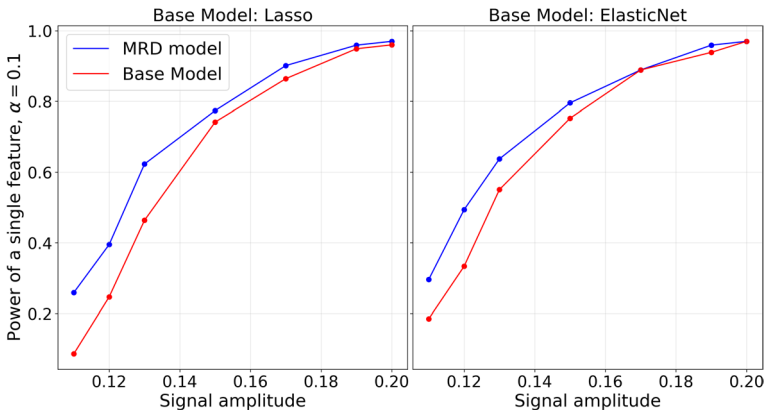


Fig. 11 Optimizing and testing for a single hypothesis on synthetic data with correlated multivariate Gaussian features and a simulated response that follows a non-linear model. The significance level is fixed and equals to $\alpha = 0.1$. The empirical power is presented as a function of the signal strength c , and evaluated by averaging over 100 independent experiments. Left: MRD lasso and lasso. Right: MRD elastic net and elastic net

G.3 Additional experiment with interaction model

In this experiment, we set the auto-correlation parameter $\rho = 0.25$ and choose the $Y|X$ to be $Y = \sum_{j=1}^{15} X_{2j}X_{2j-1} + \epsilon$ (i.e., $Y | X$ follows **M2**), such that in total there are 30 non-null elements out of total $d = 100$ features. Figure 2 presents the empirical power and FDR (target level of $q = 0.2$) obtained by a base model formulated as a neural network. We use the same network architecture and training procedure as described in Section D, but with total $n = 4000$ observations ($m = 2000$), learning-rate of 0.006, fit the model to a total of 200 epochs, and set the CancelOut regularization parameter to 0.02. (In this experiment, linear predictive models attain zero power, and hence omitted.)




Fig. 12 Q-Q plots of the p-values produced in the experiment from Table 6 compared to the uniform distribution $U[0, 1]$. Left: p-values correspond to the null features. Right: p-values correspond to the non-null features

Table 6 Synthetic experiments with correlated multivariate Gaussian features and a simulated response that follows a non-linear model with $g(X) = \sum_{j \in \mathcal{H}_1} \sin(X_j \beta_j)$

c	MRD Lasso		Lasso		% imp. of power	MRD Elastic Net		Elastic Net		% imp. of power
	Power	FDR	Power	FDR		Power	FDR	Power	FDR	
0.08	0.143	0.066	0.081	0.023	76.5	0.153	0.075	0.113	0.055	35.4
0.10	0.324	0.088	0.233	0.042	39.1	0.375	0.103	0.287	0.070	30.7
0.12	0.550	0.106	0.486	0.074	13.2	0.577	0.123	0.524	0.074	10.1
0.14	0.731	0.095	0.702	0.060	4.1	0.753	0.105	0.715	0.074	5.3
0.16	0.854	0.092	0.838	0.063	1.9	0.856	0.102	0.842	0.070	1.7
0.18	0.919	0.090	0.914	0.069	0.5	0.917	0.097	0.915	0.071	0.2

All standard errors are below 0.02. The other details are as in Table 1

G.4 Experiments with sum of sines model

In this section, we provide additional experiment with a different conditional modeling $Y | X$. Here, we use the same synthetic experiment setting as in the experiments with varying signal strength from Sect. 5.1.2, however generate $Y | X$ that follows **M3**. Table 6 presents the empirical power and FDR ($q = 0.2$) obtained when using the elastic-net and lasso as the base predictive models. Following that table, the MRD achieves a gain in power compared to the base model, and the FDR is controlled.

G.5 Experiments under model misspecification

A limitation of the model- X randomization test is that P_X must be known in order to sample the dummy features \tilde{X}_j , which are also used to formulate our RD loss. In this section,




Fig. 13 Experiment with dummy features whose quality varies with $\hat{\rho}$. Here, the true P_X follows a multivariate Gaussian distribution with with auto-correlation $\rho = 0.5$. Left: sum of covariance diagnostics over all $j = 1, \dots, p$ (lower values correspond to more accurate dummies) as a function of $\hat{\rho}$. Right: empirical power and FDR

we demonstrate that the feature selection pipeline applied with the proposed MRD method is fairly robust to unknown or estimated distributions. For ease of computations, we conduct our experiments by choosing lasso or elastic net as the base models, fitted in the same way described in Section D.

Varying the estimation quality of the conditionals

We begin with an experiment that demonstrates the effect of the density estimation on the empirical FDR. We generate features as in Sect. 5.1, but with an auto-correlation coefficient $\rho = 0.5$, higher than the one used in the synthetic experiments from that section. Here, the model of $Y | X$ follows **M2** with a fixed signal strength $c = 0.16$, where we use lasso as a base predictive model. To control the quality of the estimation of P_X , we sampled the dummy features \tilde{X}_j from $Q_{X_j|X_{-j}}$, where Q_X also follows a multivariate Gaussian however with an auto-correlation coefficient $\hat{\rho}$ that can differ from the true $\rho = 0.5$. We evaluate the quality of the dummies $\tilde{X}_j \sim Q_{X_j|X_{-j}}$ using the covariance goodness-of-fit diagnostic, presented by Romano et al, 2019, Section 5. In short, this diagnostic presents an unbiased estimate of the distance between $\text{Cov}(X_j, X_{-j})$ and $\text{Cov}(\tilde{X}_j, X_{-j})$, where larger values imply worse estimation of the conditional. The left panel of Fig. 13 depicts the sum of the covariance diagnostics over $j = 1, \dots, p$. Each point in that graph is the average value taken over 50 independent data sets, where each is of size 10, 000 samples. As can be seen, the value of the sum covariance diagnostics is close to zero when $\hat{\rho} = 0.5$, as expected, and gets higher as $\hat{\rho}$ is far from the true value of $\rho = 0.5$.

Next, we turn to study the effect of sampling inaccurate dummies on the performance of our proposed method. To this end, we use the inaccurate dummies (sampled from $Q_{X_j|X_{-j}}$) both for fitting the MRD model and for computing the test p-values. The right panel of Fig. 13 depicts the empirical power and FDR of both MRD lasso and its base version, evaluated on 50 independent data sets of size $n = 1000$ each. Following that figure, we observe that the MRD lasso is slightly more sensitive to the use of inaccurate dummy features, however a violation in the FDR control occurs only when $\hat{\rho} \geq 0.8$ for which the sum of the covariance diagnostics is above 10. This experiment illustrates that our method is fairly robust to inaccurate estimation of P_X , where invalid results may be obtained when this estimation is of poor quality. Importantly, this statement holds true for all model-X randomization tests, as the generated p-values are valid only when P_X is known. In terms of power, the MRD lasso presents, once again, a solid gain in performance compared to lasso when $\hat{\rho} = \rho$, where the FDR is controlled. (Recall that the correlation structure here is more significant than the synthetic experiments from Sect. 5.1, in which $\rho \in \{0.1, 0.25\}$.)

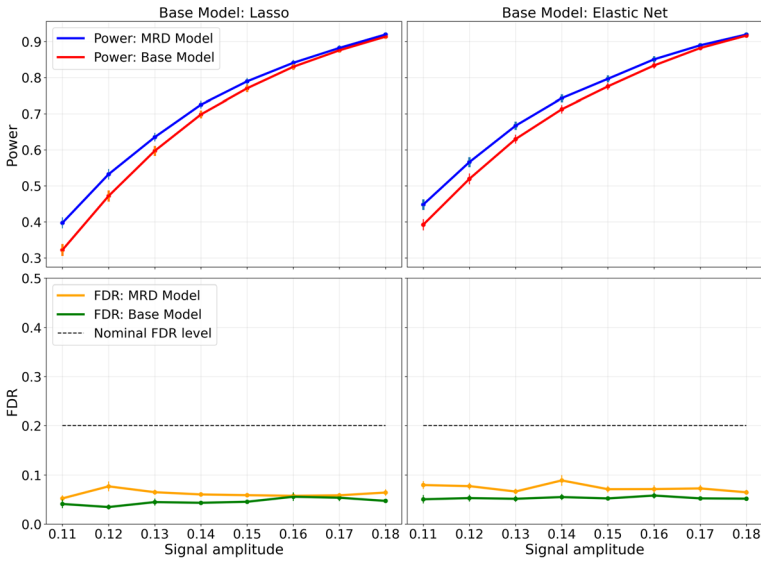


Fig. 14 Robustness experiments with correlated multivariate Gaussian features and a simulated response that follows a linear model with varying signal strength c . The mean and covariance of P_X are estimated from the data. Left: power (top) and FDR (bottom) of MRD lasso and lasso. Right: power (top) and FDR (bottom) of MRD elastic net and elastic net. The performance metrics are averaged over 100 independent experiments

Multivariate Gaussian

Consider the experimental setup from the experiments with varying signal strength in Sect. 5.1, where $X \sim \mathcal{N}(\mu, \Sigma)$ is drawn from an auto-regression model. There, we fit the MRD models and apply the HRT by drawing \tilde{X}_j from the true $P_{X_j|X_{-j}}$. Here, we conduct similar experiments, but sample \tilde{X}_j from $\hat{P}_{X_j|X_{-j}}$, where $\hat{P}_X \sim \mathcal{N}(\hat{\mu}, \hat{\Sigma})$ with mean $\hat{\mu}$ and covariance matrix $\hat{\Sigma}$ that are estimated from the observed data. Following Figs. 14, 15, we can see that all predictive models—including ours—result in FDR below the nominal 20% level, both in the linear and non-linear settings. Notice that our MRD approach consistently improves the power of each base model, which is in line with the experiments from Sect. 5.1.

Gaussian mixture model

We now present a more challenging experiment in which X is drawn independently from a multivariate Gaussian mixture model:

$$X \sim \begin{cases} \mathcal{N}(0, \Sigma_1), & \text{w.p. } \frac{1}{3}, \\ \mathcal{N}(0, \Sigma_2), & \text{w.p. } \frac{1}{3}, \\ \mathcal{N}(0, \Sigma_3), & \text{w.p. } \frac{1}{3}. \end{cases}$$

Above, Σ_1, Σ_2 and Σ_3 have the same correlation structure described in Sect. 5.1, with auto-correlation coefficients $\rho_1 = 0.1, \rho_2 = 0.2,$ and $\rho_3 = 0.3,$ respectively. Similar to the experiments presented above, we assume that X follows a multivariate Gaussian and we estimate its parameters from the observed data. This allows us to study the robustness of the proposed MRD approach to model misspecification. We follow the experimental




Fig. 15 Robustness experiments with correlated multivariate Gaussian features and a simulated response that follows a non-linear model with varying signal strength c . The other details are as in Fig. 14




Fig. 16 Robustness experiments with correlated multivariate Gaussian mixture features and a simulated response that follows a linear model with varying signal strength c . The other details are as in Fig. 14

protocol described in the experiments with varying signal strength from Sect. 5.1.2 (with **M1** for the non-linear setting), and summarize the results in Figs. 16, 17. As can be seen, the FDR obtained by all models is below the desired 20% level both in the linear and non-linear settings, indicating the robustness of the feature selection procedure. Moreover, our




Fig. 17 Robustness experiments with correlated multivariate Gaussian mixture features and a simulated response that follows a non-linear model with varying signal strength c . The other details are as in Fig. 14




Fig. 18 Experiments with correlated multivariate Student t -distribution features and a simulated response with varying signal strength c . Power and FDR are evaluated on 50 independent data sets. Left: response that follows a linear model, where the dummies are sampled from a naive multivariate Gaussian approximation of P_X . Middle: response that follows a non linear model, where the dummies are also sampled from a naive multivariate Gaussian. Right: response that follows a non linear model, where the dummies are sampled from a GMM approximation of P_X

MRD approach yields higher power than the base models, demonstrating that the proposed method is also effective when the distribution of X is unknown.

Multivariate Student’s t -distribution

In the experiments from previous paragraphs, we use a simple Gaussian approximation of P_X , and we always obtain a control of the FDR. Now, we turn to a scenario where such an approximation does not result in FDR control, and demonstrate that using a better density estimation for P_X addresses tackles this limitation. To this end, we generate $n = 1000$ samples of X , where X follows a multivariate Student’s t -distribution with zero mean, covariance matrix with entries $\Sigma_{i,j} = \rho^{|i-j|}$ where $\rho = 0.1$, and $\nu = 5$ degrees of freedom; refer to (Romano et al, 2019, Section 6.5) for the precise definition of this data generating function. We begin this experiment by sampling Y that follows a linear model (i.e., **M1**) of a varying signal strength c , as described in Sect. 5.1 of the main manuscript, where we use the same multivariate Gaussian approximation of P_X as in the previous paragraphs. The

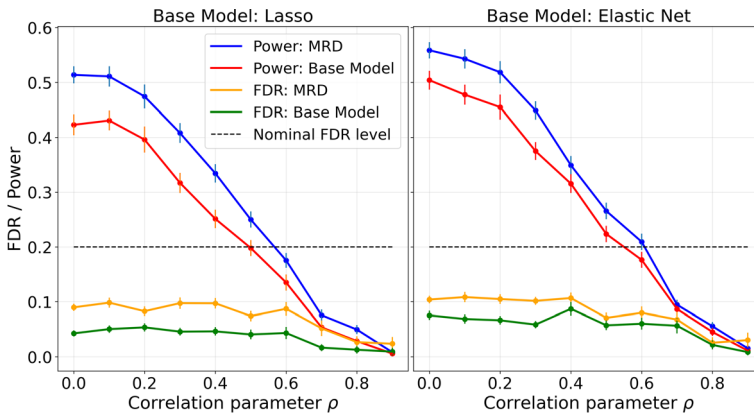


Fig. 19 Experiments with correlated features that follow $\mathcal{N}(0, \Sigma)$ where $\Sigma_{i,j} = \rho^{|i-j|}$ with varying auto-correlation parameter ρ , and a simulated response. Power and FDR are evaluated on 100 independent data sets

left panel of Fig. 18 shows the empirical power and FDR evaluated on 50 independent data sets, comparing the performance of an elastic net model and its MRD version. Next, we generate the response Y using the non linear model **M2**, and again vary the signal strength c as described in Sect. 5.1.2 of the main manuscript. As before, we approximate the distribution of X by fitting a multivariate Gaussian on the observed data. The middle panel of Fig. 18 presents the empirical power and FDR of this non linear setting, demonstrating how the MRD elastic net fails to control the FDR while the base elastic net model is below the nominal FDR rate. This violation is due to the inaccurate estimation of P_X provided by the naive multivariate Gaussian fit. To alleviate this, we now turn to approximate P_X using a Gaussian Mixture Model (GMM) (Reynolds, 2009) with 3 components, where we use Python’s `sklearn` package to fit such a GMM model. The right panel of Fig. 18 presents the empirical power and FDR obtained by sampling the dummy features from the fitted GMM model (Gimenez et al., 2019), instead of a naive multivariate Gaussian, where we use the same non-linear setting as in the middle panel. Following the right panel of that figure, we can see that the FDR is controlled across all signal strength values, showing that by using better techniques to approximate P_X we get more reliable results. This is in line with many other techniques proposed in the model-X literature that offer practical solutions to sampling more accurate dummies, as discussed in Sect. 4.1 of the main manuscript. Lastly, in terms of power, the one can see that the MRD models outperform their base versions across the board, emphasizing the superiority of our proposal.

G.6 Experiments with varying auto-correlation parameter ρ

In this section, we study the effect of the correlation between features on the proposed MRD approach. To this end, we follow the experimental setting described in Sect. 5.1, where $Y | X$ follows **M2** with a fixed signal strength $c = 0.14$ and $m = 400$, and vary the auto-correlation parameter ρ . Figure 19 presents the empirical power and FDR, evaluated over 100 independent data sets, as a function of ρ for lasso, elastic net and their MRD versions. As can be seen, the power of all methods decreases as ρ increases, indicating the challenge of testing for conditional independence with highly correlated features. However, the MRD approach results in higher power for almost every ρ while maintaining the FDR

Table 7 Synthetic experiments with correlated multivariate Gaussian features of dimension $d = 100$, and a simulated response that follows a non-linear model with varying number of samples n

n	Lasso		MRD Lasso		dCRT		Knockoffs	
	Power	FDR	Power	FDR	Power	FDR	Power	FDR
600	0.24	0.02	0.25	0.03	0.25	0.06	0.13	0.05
1000	0.50	0.02	0.52	0.01	0.52	0.05	0.42	0.05
1400	0.70	0.02	0.72	0.02	0.72	0.05	0.59	0.08

The right column represents the result of the dCRT procedure, and the rest represent the CV-HRT procedure with 20 folds. The empirical FDR (nominal level $q = 0.1$) and power are evaluated by averaging over 50 independent experiments. All standard errors are below 0.02

under control, demonstrating the advantage of the MRD approach even when working with correlated features.

G.7 A comparison of CV-HRT to dCRT and Knockoffs

In this section, we use simulated data to compare the performance of CV-HRT (Tansey et al., 2021, Algorithm 4) to that of dCRT (Liu et al., 2020) and of the model-X knockoffs (Candès et al., 2018). (Recall that the former relies on cross-validating and thus avoids the naive data-splitting of the vanilla HRT procedure.) To this end, we generate $X \in \mathbb{R}^d$ that follows an autoregression model with correlation coefficient $\rho = 0.6$, and set $d = 100$. Then, we generate a response variable $Y = (X^T \beta + \epsilon)^3$, where $\epsilon \sim \mathcal{N}(0, 1)$ is a noise component, and $\beta \in \mathbb{R}^d$ is a sparse vector whose 30 first entries are equal to $c = 1.5$; we set the rest 70 elements in β to zero. We choose the elastic net to be the base predictive model for both CV-HRT and its MRD version, and compare the performance of these two tests to dCRT and to the model-X knockoffs. The CV-HRT is implemented by setting the number of folds to be equal to $K = 20$. We implemented the dCRT in Python, by following the official R package provided by the authors⁴. we run the dCRT method using the default choice of parameters, provided in the official R package. We run the model-X knockoffs with a second order estimation for generating the knockoffs, by using the package provided by Romano et al. (2019) with its default choice of parameters.⁵

Table 7 compares the performance of the above four frameworks (averaged over 50 trials) for a varying number of samples n , where the target FDR level is fixed and equals to $q = 0.1$. Following that table, we can see that the gap between CV-HRT and the dCRT is relatively small, and the MRD approach closes this gap even further, achieving essentially the same power as dCRT. The model-X knockoffs achieves lower power than all other frameworks. In all experiments, the empirical FDR is controlled, where the MRD version of CV-HRT tends to have a lower empirical FDR compared to dCRT and the model-X knockoffs. Lastly, we note that CV-HRT requires fitting K different models (20 in this experiment), less than the d leave-one-covariate-out models that are required to implement the dCRT test. Therefore, in general, the computational complexity of CV-HRT may be lower than that of the dCRT.

⁴ <https://github.com/moleibobliu/Distillation-CRT>.

⁵ We use the software package from <https://github.com/mnesia/deepknockoffs>.




Fig. 20 Semi-synthetic experiments with real HIV mutation features and a simulated response that follows a non-linear model with varying signal strength c . Empirical power and FDR with $q = 0.2$ are evaluated over 20 random train/test splits of the data. The other details are as in Fig. 3

Appendix H Supplementary details on real data experiments

H.1 Description of the data

In Sect. 5.2 of the main manuscript, we analyze which genetic variations are linked to changes in human immunodeficiency virus drug-resistance. We download the data set from http://hivdb.stanford.edu/pages/published_analysis/genophenoPNAS2006 and deploy the same pre-processing suggested in (Romano et al., 2019). Specifically, we selected $d = 150$ features in total, where 75 of them are chosen because they are already reported as important,⁶ and the rest 75 are those who have the most frequently occurring mutations. After discarding samples that have missing values, we have 1555 examples with 150 features each.

H.2 Semi-synthetic experiments with a non-linear model

Here we provide the complementary results of the semi-synthetic non-linear data experiments, discussed in Sect. 5.2. Figure 20 presents the empirical power and FDR evaluated over 20 realizations of the data set (random train/test splits of the data), using lasso, elastic net and neural network as the base models. In addition to the discussion in Sect. 5.2, it can be seen that the improvement of the neural network is consistent but less significant than those of the linear predictive models (lasso and elastic net). However, although we misspecify the distribution P_X , the empirical FDR is below the nominal level $q = 0.2$ in the neural network as well as lasso and elastic net.

H.3 Semi-synthetic experiments with a linear model

In addition to the non-linear semi-synthetic experiments described in Sect. 5.2 of the main manuscript, in this section we perform variable selection with real human immunodeficiency virus mutation features and a simulated response that follows a *linear* model with varying signal strength c . That is, we simulate Y as described in (6) such that $y | X$ follows **M1**, while treating the real features X as fixed. The results are presented in Fig. 21. As demonstrated, the FDR obtained by all methods is below the target level $q = 0.2$, indicating the

⁶ We use the database from <https://hivdb.stanford.edu/dr-summary/comments/PI>.




Fig. 21 Semi-synthetic experiments with real human immunodeficiency virus mutation features and a simulated response that follows a linear model with varying signal strength c . Empirical power and FDR with $q = 0.2$ are evaluated over 50 random train/test splits of the data. The other details are as in Fig. 3

Table 8 Semi-Synthetic experiments with real features and a simulated response that follows a linear model with varying number of samples n and a fixed signal strength $c = 3$

n	MRD Lasso		Lasso		% imp. of power
	Power	FDR	Power	FDR	
100	0.065	0.055	0.040	0.026	62.5
150	0.223	0.062	0.172	0.026	29.7
200	0.341	0.081	0.290	0.060	17.6
250	0.462	0.085	0.430	0.041	7.4
300	0.579	0.088	0.548	0.050	5.7
350	0.646	0.078	0.629	0.050	2.7
500	0.820	0.080	0.810	0.054	1.2
600	0.897	0.075	0.900	0.049	-0.3

The empirical FDR (nominal level $q = 0.2$) and power are evaluated by averaging over 50 independent experiments. All standard errors are below 0.02

robustness of the selection procedure to model misspecification. (Recall that we approximate the unknown P_X as a multivariate Gaussian whose parameters are estimated from the whole data.) In terms of power, our MRD methods outperform the base predictive models. This is in line with the overall tendency of the experiments presented in this paper, and, in particular, with the ones the correspond to the non-linear semi-synthetic case.

H.4 Semi-synthetic experiments with a varying number of samples

Herein, we extend the experiments from the previous sub-section and study the performance of our method as a function of the sample size. We follow the linear semi-synthetic setting described in Section H.3 with a fixed signal strength $c = 3$, however sample data sets of different size by randomly choosing a subset of $n < 1555$ observations from the whole data set. We apply the controlled variable selection procedure only on the selected

subset, where the marginal distribution of P_X is estimated using the entire data set; see section 5.2. The results obtained when choosing lasso to be the base predictive model are summarized in Table 8. As can be seen, the FDR is empirically controlled across all settings, and the power increases with the sample size. Also, consistent with previous experiments, the MRD approach outperforms the baseline method in terms of power.

Acknowledgements We sincerely thank Matteo Sesia for insightful comments about an earlier manuscript draft. S.S and Y.R. were supported by the Israel Science Foundation (Grant 729/21). Y.R. also thanks the Career Advancement Fellowship, Technion, for providing research support.

Author Contributions All authors contributed to the study conception and design. All authors participated in writing the manuscript.

Funding S.S and Y.R. were supported by the Israel Science Foundation (Grant 729/21). Y.R. was also supported by the Career Advancement Fellowship, Technion.

Availability of data and materials All supplementary material and data are attached to the submission.

Code availability Code for reproducing our main results is attached to the submission.

Declarations

Conflict of interest The authors declare that they have no conflict of interests.

References

- Arjovsky, M., Bottou, L., Gulrajani, I., et al. (2019). Invariant risk minimization. arXiv preprint [arXiv:1907.02893](https://arxiv.org/abs/1907.02893).
- Barber, R. F., & Candès, E. J. (2015). Controlling the false discovery rate via knockoffs. *Annals of Statistics*, 43(5), 2055–2085.
- Bates, S., Sesia, M., Sabatti, C., et al. (2020). Causal inference in genetic trio studies. *Proceedings of the National Academy of Sciences*, 117(39), 24,117–24,126.
- Bellot, A., & van der Schaar, M. (2019). Conditional independence testing using generative adversarial networks. In: *Advances in neural information processing systems* 32.
- Bender, R., & Lange, S. (2001). Adjusting for multiple testing—when and how? *Journal of Clinical Epidemiology*, 54(4), 343–349.
- Benjamini, Y., & Hochberg, Y. (1995). Controlling the false discovery rate: a practical and powerful approach to multiple testing. *Journal of the Royal Statistical Society: Series B (Methodological)*, 57(1), 289–300.
- Benjamini, Y., & Yekutieli, D. (2001). The control of the false discovery rate in multiple testing under dependency. *Annals of statistics* (pp. 1165–1188).
- Benner, C., Spencer, C. C., Havulinna, A. S., et al. (2016). Finemap: Efficient variable selection using summary data from genome-wide association studies. *Bioinformatics*, 32(10), 1493–1501.
- Borisov, V., Haug, J., & Kasneci, G. (2019). CancelOut: A layer for feature selection in deep neural networks. In *Artificial neural networks and machine learning - ICANN*.
- Boyd, S., Parikh, N., Chu, E., et al. (2011). Distributed optimization and statistical learning via the alternating direction method of multipliers. *Foundations and Trends in Machine Learning*, 3(1), 1–122.
- Burns, C., Thomason, J., & Tansey, W. (2020). Interpreting black box models via hypothesis testing. In: *Proceedings of the 2020 ACM-IMS on foundations of data science conference* (pp. 47–57).
- Candès, E., Fan, Y., Janson, L., et al. (2018). Panning for gold: Model-X knockoffs for high dimensional controlled variable selection. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 80(3), 551–577.
- Chen, J., & Li, H. (2013). Variable selection for sparse dirichlet-multinomial regression with an application to microbiome data analysis. *The Annals of Applied Statistics*, 7(1).
- Chen, J., Song, L., & Wainwright, M., et al. (2018). Learning to explain: An information-theoretic perspective on model interpretation. In *International conference on machine learning*, PMLR (pp. 883–892).

- Chwialkowski, K., Strathmann, H., & Gretton, A. (2016). A kernel test of goodness of fit. In *International conference on machine learning*, PMLR (pp. 2606–2615).
- Gimenez, J. R., Ghorbani, A., & Zou, J. (2019). Knockoffs for the mass: new feature importance statistics with false discovery guarantees. In *The 22nd international conference on artificial intelligence and statistics*, PMLR (pp. 2125–2133).
- Gretton, A., Sriperumbudur, B. K., & Sejdinovic, D., et al. (2012). Optimal kernel choice for large-scale two-sample tests. In *Advances in Neural Information Processing Systems* 25, 2012.
- Hawinkel, S., Mattiello, F., Bijens, L., et al. (2019). A broken promise: Microbiome differential abundance methods do not control the false discovery rate. *Briefings in Bioinformatics*, 20(1), 210–221.
- Katsevich, E., & Ramdas, A. (2020). A theoretical treatment of conditional independence testing under model-X. arXiv preprint [arXiv:2005.05506](https://arxiv.org/abs/2005.05506).
- Khani, F., & Liang, P. (2021). Removing spurious features can hurt accuracy and affect groups disproportionately. In *Proceedings of the 2021 ACM conference on fairness, accountability, and transparency* (pp. 196–205).
- Kingma, D. P., & Ba, J. (2014). Adam: A method for stochastic optimization. arXiv preprint [arXiv:1412.6980](https://arxiv.org/abs/1412.6980).
- König, G., Molnar, C., & Bischl, B., et al. (2021). Relative feature importance. In: *25th International conference on pattern recognition*, IEEE (pp. 9318–9325).
- Lemhadri, I., Ruan, F., & Tibshirani, R. (2021). LassoNet: Neural networks with feature sparsity. In *International conference on artificial intelligence and statistics*, PMLR (pp. 10–18).
- Liu, M., Katsevich, E., & Janson, L., et al. (2020). Fast and powerful conditional randomization testing via distillation. arXiv preprint [arXiv:2006.03980](https://arxiv.org/abs/2006.03980).
- Liu, Q., Lee, J., & Jordan, M. (2016). A kernelized Stein discrepancy for goodness-of-fit tests. In *International conference on machine learning*, PMLR (pp. 276–284).
- Lu, Y. Y., Fan, Y., & Lv, J., et al. (2018). DeepPINK: Reproducible feature selection in deep neural networks. In *Advances in Neural Information Processing Systems* 31.
- Manolio, T. A., Collins, F. S., Cox, N. J., et al. (2009). Finding the missing heritability of complex diseases. *Nature*, 461(7265), 747–753.
- McCoy, R. T., Pavlick, E., & Linzen, T. (2019). Right for the wrong reasons: Diagnosing syntactic heuristics in natural language inference. arXiv preprint [arXiv:1902.01007](https://arxiv.org/abs/1902.01007).
- Parikh, N., & Boyd, S. (2014). Proximal algorithms. *Foundations and Trends in Optimization*, 1(3), 127–239.
- Peters, J., Buhlmann, P., & Meinshausen, N. (2015). Causal inference using invariant prediction: identification and confidence intervals. [arxiv. Methodology](https://arxiv.org/abs/1505.04747).
- Reynolds, D. (2009). *Gaussian mixture models* (pp. 659–663). Boston: Springer.
- Rhee, S. Y., Taylor, J., Wadhera, G., et al. (2006). Genotypic predictors of human immunodeficiency virus type 1 drug resistance. *Proceedings of the National Academy of Sciences*, 103(46), 17,355–17,360.
- Romano, Y., Bates, S., & Candès, E. J. (2020). Achieving equalized odds by resampling sensitive attributes. In: *Advances in Neural Information Processing Systems* 33.
- Romano, Y., Sesia, M., & Candès, E. (2019). Deep knockoffs. *Journal of the American Statistical Association*, 115(532), 1861–1872.
- Sagawa, S., Koh, P. W., & Hashimoto, T. B., et al. (2019). Distributionally robust neural networks for group shifts: On the importance of regularization for worst-case generalization. arXiv preprint [arXiv:1911.08731](https://arxiv.org/abs/1911.08731).
- Sanjabi, M., Ba, J., & Razaviyayn, M., et al. (2018). Solving approximate wasserstein gans to stationarity. arXiv preprint [arXiv:1802.08249](https://arxiv.org/abs/1802.08249).
- Sesia, M., Bates, S., & Candès, E., et al. (2020a). Controlling the false discovery rate in GWAS with population structure. bioRxiv.
- Sesia, M., Katsevich, E., Bates, S., et al. (2020). Multi-resolution localization of causal variants across the genome. *Nature Communications*, 11(1), 1–10.
- Sesia, M., Sabatti, C., & Candès, E. J. (2019). Gene hunting with hidden Markov model knockoffs. *Biometrika*, 106(1), 1–18.
- Shah, R. D., & Peters, J. (2020). The hardness of conditional independence testing and the generalised covariance measure. *The Annals of Statistics*, 48(3), 1514–1538.
- Srivastava, N., Hinton, G., Krizhevsky, A., et al. (2014). Dropout: A simple way to prevent neural networks from overfitting. *The Journal of Machine Learning Research*, 15(1), 1929–1958.
- Sudarshan, M., Puli, A., & Subramanian, L., et al. (2021). Contra: Contrarian statistics for controlled variable selection. In *International conference on artificial intelligence and statistics*, PMLR (pp. 1900–1908).

- Sudarshan, M., Tansey, W., & Ranganath, R. (2020). Deep direct likelihood knockoffs. *Advances in Neural Information Processing Systems* 33.
- Tansey, W., Veitch, V., & Zhang, H., et al. (2021). The holdout randomization test for feature selection in black box models. *Journal of Computational and Graphical Statistics* 1–37.
- Yamada, Y., Lindenbaum, O., & Negahban, S., et al. (2020). Feature selection using stochastic gates. In *International conference on machine learning*, PMLR (pp. 10,648–10,659).
- Yoon, J., Jordon, J., & van der Schaar, M. (2019). INVASE: Instance-wise variable selection using neural networks. In *International conference on learning representations*.
- Yu, Y. (2013). On decomposing the proximal map. In *Advances in Neural Information Processing Systems* 26.
- Zhang, K., Peters, J., & Janzing, D., et al. (2012). Kernel-based conditional independence test and application in causal discovery. arXiv preprint [arXiv:1202.3775](https://arxiv.org/abs/1202.3775).
- Zou, H., & Hastie, T. (2005). Regularization and variable selection via the elastic net. *Journal of the Royal Statistical Society, Series B*, 67, 301–320.

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

Springer Nature or its licensor (e.g. a society or other partner) holds exclusive rights to this article under a publishing agreement with the author(s) or other rightsholder(s); author self-archiving of the accepted manuscript version of this article is solely governed by the terms of such publishing agreement and applicable law.