
More Powerful Selective Kernel Tests for Feature Selection

Jen Ning Lim
University College London

Makoto Yamada
Kyoto University, RIKEN AIP

Wittawat Jitkrittum
MPI for Intelligent Systems, Tübingen

Yoshikazu Terada
Osaka University, RIKEN AIP

Shigeyuki Matsui
Nagoya University

Hidetoshi Shimodaira
Kyoto University, RIKEN AIP

Abstract

Refining one’s hypotheses in the light of data is a common scientific practice; however, the dependency on the data introduces selection bias and can lead to specious statistical analysis. An approach for addressing this is via conditioning on the selection procedure to account for how we have used the data to generate our hypotheses, and prevent information to be used again after selection. Many selective inference (a.k.a. post-selection inference) algorithms typically take this approach but will “over-condition” for sake of tractability. While this practice yields well calibrated statistic tests with controlled false positive rates (FPR), it can incur a major loss in power. In our work, we extend two recent proposals for selecting features using the Maximum Mean Discrepancy and Hilbert Schmidt Independence Criterion to condition on the minimal conditioning event. We show how recent advances in multiscale bootstrap makes conditioning on the minimal selection event possible and demonstrate our proposal over a range of synthetic and real world experiments. Our results show that our proposed test is indeed more powerful in most scenarios.

1 INTRODUCTION

Most statistical methods implicitly assume that parameters of the statistical investigation are fixed a priori; that is, the choice of model, hypothesis to test,

and parameters to be estimated do not change before the data is inspected. Failure to satisfy this can lead to disturbing properties such as uncalibrated p -values (Simmons et al., 2011; Gelman and Loken, 2013). The field of selective inference (SI) considers a modernised version of statistical analysis where we first explore the data and determine relevant parameters for our investigation. Then, SI aims to provide valid inference under the model chosen by the data (Fithian et al., 2014). In our work, we extend two algorithms that first select a set of features then perform hypothesis testing on each of the selected features to determine whether it is statistically significant.

One of the approaches in the field of SI is conditioning on how the data has been used during the initial selection phase (Fithian et al. (2014); Lee et al. (2016); Fithian et al. (2015)). This approach may be difficult to use since it requires an explicit characterisation of the selection procedure and the conditional distribution of the test statistic, both of which can be difficult to obtain. Fortunately, one of the key developments that has allowed many SI algorithms to be tractable is the polyhedral lemma (Lee et al., 2016; Tibshirani et al., 2016). Assume that the test statistic is normally distributed before the selection. The polyhedral lemma states that if the selection event can be written as a set of linear constraints, then its conditional post-selective distribution follows a truncated normal distribution (Lee et al., 2016, Theorem 5.2). This result has been successfully applied to non-parametric kernel methods for selecting informative features using the Hilbert Schmidt Independence Criterion (Yamada et al., 2018), and the Maximum Mean Discrepancy (Yamada et al., 2019), as well as multiple model comparison (Lim et al., 2019).

A subtlety with SI is that the power of selective hypothesis tests (i.e., tests with null hypothesis that is determined by data and so *random*) depends upon our choice of what to condition on. If we condition on too little, the test will have uncontrolled false positive rate.

Proceedings of the 23rd International Conference on Artificial Intelligence and Statistics (AISTATS) 2020, Palermo, Italy. PMLR: Volume 108. Copyright 2020 by the author(s).

If we condition on too much, it can incur a loss of power (Fithian et al., 2014). This observation has driven research efforts to curate more powerful hypothesis tests that have higher “left-over” information. These proposals include randomising the data used during selection (Tian et al., 2018); a careful characterisation of how the data has been used during the selection process as there are different costs for variable selection and target formulation (Liu et al., 2018); and conditioning on the minimal set, i.e., condition only on what is necessary for the test to be valid but not more (Liu et al., 2018; Terada and Shimodaira, 2019). The last idea forms the basis of what we propose in this work.

In the present work, we consider the problem of selecting a subset of informative features with selective inference. We study two related problem settings. In the first setting (Section 3), given two samples, the goal is to select a subset of features for which the marginal distributions (restricted to the selected subset) of the two underlying distributions significantly differ. In the second setting (Section 4), given a joint sample of covariate and response variables, the goal is to select a subset of covariate variables whose dependency on the response is statistically significant. While the selective tests of Yamada et al. (2019, 2018) are applicable to these problems, and have a tractable null distribution due to the use of the polyhedral lemma, these tests do not consider the minimal condition set, meaning that the tests may be overly conservative.

We generalize the tests of Yamada et al. (2019) and Yamada et al. (2018), for the two settings respectively, and propose tests that condition on the minimal conditioning set by using the selective multiscale bootstrap (Terada and Shimodaira, 2017, 2019). For the second problem, we further propose a new estimator for the Hilbert Schmidt Independence Criterion (HSIC) that takes the form of an incomplete U-statistic. We show that the new estimator leads to a test that has higher power than the test of Yamada et al. (2018) which relies on the block estimator. In experiments (Section 5) on both synthetic and real problems, we show that the new tests have well-controlled false positive rate, and are more powerful than their respective original tests when the number of features is large, and the number of selected features is larger than one.

2 BACKGROUND

In this section, we review the Maximum Mean Discrepancy (MMD) and Hilbert Schmidt Independence Criterion (HSIC) which are used as our criteria to select features as well as briefly introduce the concept of multiscale bootstrap. In Section 3, we use MMD to select features that have significantly different marginal

distributions and in Section 4, we use HSIC to select features which have a significant dependence on the response variable respectively.

Maximum Mean Discrepancy (MMD) For a distribution P and a positive definite kernel K , the mean embedding of P is defined as $\mu_P = \mathbb{E}_{x \sim P}[K(\cdot, x)]$ (Smola et al., 2007). The Maximum Mean Discrepancy (MMD) is a pseudo metric between two distributions P and Q and is defined as $\text{MMD}(P, Q) = \|\mu_P - \mu_Q\|_K$, where $\|\cdot\|_K$ denotes the norm in the reproducing kernel Hilbert space (RKHS) associated with K . If K is a characteristic kernel, then $\text{MMD}^2(P, Q) = 0 \iff P = Q$ (Gretton et al., 2012). An example of a characteristic kernel is the Gaussian kernel. It can be shown that the squared MMD can be written equivalently as $\text{MMD}^2(P, Q) = \mathbb{E}_{z, z' \sim P \times Q}[h(z, z')]$ where $z := (x, y)$ and $h(z, z') := K(x, x') + K(y, y') - K(x', y) - K(x, y')$. Given samples $\mathbf{x} := \{x_i\}_{i=1}^n$ and $\mathbf{y} := \{y_i\}_{i=1}^n$ of size n as i.i.d. draws from P and Q respectively, an unbiased estimator is the U-statistic $\widehat{\text{MMD}}_u^2(\mathbf{x}, \mathbf{y}) = \frac{1}{n(n-1)} \sum_{i \neq j} h(z_i, z_j)$. In our work, we focus on a parametric bootstrap resampling procedure for multiscale bootstrap and thus we use estimators with normal asymptotic distributions such as the linear-time estimator $\widehat{\text{MMD}}_l^2 = \frac{2}{n} \sum_{i=1}^{n/2} h(z_{2i}, z_{2i-1})$ (Gretton et al., 2012), and the incomplete U-statistic estimator (Blom, 1976; Janson, 1984) proposed by Yamada et al. (2019): $\widehat{\text{MMD}}_{Inc}^2 = \frac{1}{|\mathcal{D}_n|} \sum_{(i,j) \in \mathcal{D}_n} h(z_i, z_j)$ where \mathcal{D}_n is random and sampled with replacement from $\{(i, j)\}_{i \neq j}$. Under weak assumptions, both $\widehat{\text{MMD}}_l^2$ and $\widehat{\text{MMD}}_{Inc}^2$ are asymptotically normal for both when $P = Q$ and $P \neq Q$ (Gretton et al., 2012; Yamada et al., 2019).

Hilbert Schmidt Independence Criterion (HSIC) Let $x \in \mathcal{X}$ and $y \in \mathcal{Y}$ be two random variables with joint distribution P_{xy} and their respective marginals P and Q . Let $K_{\mathcal{X}}$ and $K_{\mathcal{Y}}$ be two real-valued kernel functions defined on $\mathcal{X} \times \mathcal{X}$ and $\mathcal{Y} \times \mathcal{Y}$ respectively. The Hilbert Schmidt Independence Criterion (Gretton et al., 2005) is defined as the Hilbert-Schmidt norm of the covariance operator $\text{HSIC}(P_{xy}) = \|\mu_{xy} - \mu_P \otimes \mu_Q\|_{\text{HS}}^2$ where \otimes denotes the tensor product and $\mu_{xy} := \mathbb{E}_{(x,y) \sim P_{xy}}[K_{\mathcal{X}}(x, \cdot) \otimes K_{\mathcal{Y}}(y, \cdot)]$. The norm $\|\cdot\|_{\text{HS}}$ is induced by the inner product of the space of linear operators (that are Hilbert Schmidt). See Gretton et al. (2005, Section 2) for details. If the product kernel $K_{\mathcal{X}}K_{\mathcal{Y}}$ is characteristic on the joint domain $\mathcal{X} \times \mathcal{Y}$, then $\text{HSIC}(P_{xy}) = 0 \iff x \perp y$ (x and y are independent) (Fukumizu et al., 2008, Theorem 3). An example of such a kernel can be constructed by letting $K_{\mathcal{X}}$ and $K_{\mathcal{Y}}$ be Gaussian kernels on $\mathcal{X} \subseteq \mathbb{R}^{d_x}$ and $\mathcal{Y} \subseteq \mathbb{R}^{d_y}$ respectively. Given $\mathbf{z} := \{(x_i, y_i)\}_{i=1}^n$ consisting of n i.i.d. samples from

P_{xy} , an unbiased estimator can be computed as a U-statistic $\widehat{\text{HSIC}}_u(\mathbf{z}) = \frac{(n-4)!}{n!} \sum_{(i,j,q,r) \in \mathbf{i}_4^n} h(i,j,q,r)$ (Hoeffding, 1992; Song et al., 2012) where \mathbf{i}_4^n is the set of all 4-tuples with each index occurring only once, $h(i,j,q,r) = \frac{1}{4!} \sum_{(s,t,u,v) \in \mathcal{P}(i,j,q,r)} \mathbf{K}_{st}[\mathbf{L}_{st} + \mathbf{L}_{uv} - 2\mathbf{L}_{su}]$ is the U-statistic kernel with the sum being over $4!$ quadruples (s,t,u,v) as permutations of (i,j,q,r) , and $\mathbf{K}, \mathbf{L} \in \mathbf{R}^{n,n}$ contain entries $\mathbf{K}_{ij} = K_{\mathcal{X}}(x_i, x_j)$, $\mathbf{L}_{ij} = K_{\mathcal{Y}}(y_i, y_j)$. If x and y are dependent then $\widehat{\text{HSIC}}_u$ is asymptotically normally distributed (Song et al., 2012, Theorem 5). However, if they are independent then U-statistic is degenerate and the asymptotic distribution of $\widehat{\text{HSIC}}_u$ deviates from normal (Gretton et al., 2008; Serfling, 2009). For a given block size B , the block estimator of Zhang et al. (2018) is defined as $\widehat{\text{HSIC}}_{\text{Blo}}(\mathbf{z}) = \frac{B}{n} \sum_{i=1}^{\frac{n}{B}} \widehat{\text{HSIC}}_u(\{z_j\}_{j=(i-1)B+1}^{iB})$. If $\lim_{n,B \rightarrow \infty} n/B = \infty$, it can be shown that the block estimator is normally distributed asymptotically even when x is independent of y (Zhang et al., 2018, Section 3.2).

Multiscale Bootstrap A procedure that calculates ‘‘approximately unbiased’’ p -values is called multiscale bootstrap proposed by Shimodaira (2002); Shimodaira et al. (2004). It was initially proposed for a general statistical problem, called the problem of regions (Efron et al., 1998), where we want to compute asymptotically accurate p -values for the null hypothesis $H_0 : \mu \in H$ where H is represented by a region with $H \subseteq \mathbb{R}^d$ (called ‘‘hypothesis region’’). Efron et al. (1996) studied this problem under the normal model $y \sim \mathcal{N}(\mu, I)$ and argued that the bootstrap probabilities $\text{BP}(H) := \mathbb{P}(y \in H)$ are biased frequentist confidence measures. Furthermore, they showed that geometric quantities play a crucial role and bias corrected p -values can be produced by using the pivotal quantity $\beta_0(y) - \beta_1 \sim \mathcal{N}(0, 1)$ where $\beta_0(y)$ is the signed distance from y to ∂H (i.e., the boundary surface of H), and β_1 is the mean curvature of ∂H . More specifically, a second-order asymptotically accurate p -value is expressed as $p(H|y) := \bar{\Phi}(\beta_0(y) - \beta_1)$, i.e., $\forall \mu \in \partial H$, $\mathbb{P}(p(H|y) < \alpha) = \alpha + \mathcal{O}(n^{-1})$ (Efron et al., 1996; Shimodaira, 2002). Note that $\Phi(\cdot)$ is the CDF of the standard normal distribution and $\bar{\Phi}(x) := 1 - \Phi(x)$. However, typically $\beta_0(y)$ and β_1 are hard to determine due to either the intractability of the space or the lack of an explicit formulation in the region. Multiscale bootstrap addresses this problem with additional computation and only requires the regions to be represented by a function that indicates if $y \in H$ or $y \notin H$.

Let $\mathcal{X}_n = \{x_i\}_{i=1}^n$ be a dataset of sample size n with each element $x_i \in \mathbb{R}^d$. We assume that there is some transformation f_n such that the observed value $f_n(\mathcal{X}_n)$ follows a multivariate normal distribu-

tion, i.e., $y := f_n(\mathcal{X}_n) \sim \mathcal{N}(\mu, I)$. Typically, f_n has a factor \sqrt{n} for scaling the covariance. The main idea of multiscale bootstrap is, instead of n elements, it resamples n' elements from \mathcal{X}_n with replacement to generate $\mathcal{X}_{n'}$, then $y^* := f_n(\mathcal{X}_{n'}) \sim \mathcal{N}(y, \gamma^2 I)$ with $\gamma^2 = n/n'$, from which we estimate the desired geometric quantities $\beta_0(y)$ and β_1 using the scaling law of bootstrap probabilities (Shimodaira, 2002, 2014). It can be shown that the bootstrap probability of the region H is expressed as $\text{BP}_{\gamma^2}(H) := \mathbb{P}(y^* \in H) \approx \bar{\Phi}(\gamma^{-1}\beta_0(y) + \gamma\beta_1)$. Shimodaira (2008, 2014) proposed the normalised bootstrap z -value as $\psi_{\gamma^2}(y|H) := \gamma\bar{\Phi}^{-1}(\text{BP}_{\gamma^2}(H)) \approx \beta_0(y) + \gamma^2\beta_1$ from which p -values proposed by Efron et al. (1998), namely $p(H|y) = \bar{\Phi}(\beta_0(y) - \beta_1)$, can be calculated when $\gamma^2 = -1$, i.e., we have $\bar{\Phi}(\psi_{-1}(y|H)) = \bar{\Phi}(\beta_0(y) - \beta_1) = p(y|H)$. However, it is impossible to simulate the case where $\gamma^2 = -1$. Multiscale bootstrap tackles this problem by using a number of different sample sizes $n' \in \mathcal{M} \subset \mathbb{N}^+$. For each n' , we run bootstrap resampling of $\mathcal{X}_{n'}$ from \mathcal{X}_n for calculating the normalised bootstrap z -value $\psi_{\gamma_{n'}^2}(y|H)$ with $\gamma_{n'}^2 := n/n'$. The tuple $\{(\gamma_{n'}^2, \psi_{\gamma_{n'}^2}(y|H))\}_{n' \in \mathcal{M}}$ is used to fit a regression model $\varphi_H(\gamma^2)$ which can then be extrapolated to $\varphi_H(-1)$. The regression model $\varphi_H(\cdot)$ can be used to calculate our p -values since $p(y|H) = \bar{\Phi}(\varphi_H(-1))$. See Shimodaira and Terada (2019, Section 6.5) for several possibilities of regression models.

Selective Multiscale Bootstrap Multiscale bootstrap can be extended to the problem of selective inference, where the hypothesis is random and chosen from the data (Shimodaira and Terada, 2019; Terada and Shimodaira, 2017, 2019). In this problem, there is an additional region S called the ‘‘selective region’’ that determines the null hypothesis we are going to test. If $y \in S$ then we test $H_0 : y \in H$. However, if $y \notin S$, we ignore H and no decision is made. Terada and Shimodaira (2017, 2019) proposed the following selective p -value

$$p(H|y, S) = \frac{\bar{\Phi}(\beta_0(y|H) - \beta_1)}{\bar{\Phi}(\beta_0(y|H) - \beta_1 + \beta_0(y|S))},$$

where β_1 is the mean curvature of ∂H , and $\beta_0(y|H)$ and $\beta_0(y|S)$ are the signed distances from y to ∂H and ∂S , respectively. Under certain assumptions, its null distribution is uniform over $(0, 1)$ (Terada and Shimodaira, 2017), i.e., we have

$$p(H|y, S) | y \in S, \mu \in \partial H \sim \mathcal{U}(0, 1).$$

The calculation of $p(H|y, S)$ may again be non-trivial. The difficulty arises from the calculation of the signed distance $\beta_0(y|S)$ for the selective region S but fortunately, we can apply the same idea as non-selective multiscale bootstrap as mentioned in previous subsection. In this case, another regression model $\varphi_S(\gamma^2)$

is fitted with the bootstrap probabilities $\text{BP}_{\gamma^2}(S)$ for the region S . It can be seen that the signed distance $\beta_0(y|S)$ can then be obtained by extrapolating the model $\varphi_S(\gamma^2)$ to $\gamma^2 = 0$. In other words, the selective p -value $p(H|y, S)$ can be calculated using two regression models $\varphi_H(\cdot)$ and $\varphi_S(\cdot)$ as follows,

$$p(H|y, S) = \frac{\bar{\Phi}(\varphi_H(-1))}{\bar{\Phi}(\varphi_H(-1) + \varphi_S(0))}.$$

Algorithm 1 describes selective multiscale bootstrap

Algorithm 1 Selective Multiscale Bootstrap

```

1: procedure SELECTIVEMS( $\mathbb{1}_H, \mathbb{1}_S, \alpha, \boldsymbol{\mu}, \boldsymbol{\Sigma}$ )
2:   for  $n' \in \mathcal{M}$  do
3:      $\gamma_{n'}^2 \leftarrow \frac{\alpha}{n'}$ 
4:     Sample  $\{\mathbf{y}_i^*\}_{i=1}^B \stackrel{i.i.d.}{\sim} \mathcal{N}(\boldsymbol{\mu}, \gamma_{n'}^2 \boldsymbol{\Sigma})$ 
5:      $\text{BP}_{\gamma_{n'}^2}(H) \leftarrow \sum_{i=1}^B \mathbb{1}_H(\mathbf{y}_i^*)/B$ 
6:      $\text{BP}_{\gamma_{n'}^2}(S) \leftarrow \sum_{i=1}^B \mathbb{1}_S(\mathbf{y}_i^*)/B$ 
7:   end for
8:   Fit a model  $\varphi_H(\gamma^2)$  such that  $\varphi_H(\gamma^2) = \gamma \bar{\Phi}^{-1}(\text{BP}_{\gamma^2}(H))$ .
9:   Fit a model  $\varphi_S(\gamma^2)$  such that  $\varphi_S(\gamma^2) = \gamma \bar{\Phi}^{-1}(\text{BP}_{\gamma^2}(S))$ .
10: return  $\bar{\Phi}(\varphi_H(-1))/\bar{\Phi}(\varphi_H(-1) + \varphi_S(0))$ 
11: end procedure
    
```

algorithm where $\mathbb{1}_H(y)$ is an indicator function: it is 1 if $y \in H$, otherwise 0.

3 SELECTIVE INFERENCE WITH MMD

In this section, we propose our first test. We are concerned with the following problem,

Problem 1. *Given two distributions P and Q with common support on \mathcal{X}^d , we have n i.i.d. samples denoted as $\mathbf{X}_n = [\mathbf{x}_1, \dots, \mathbf{x}_n]^\top \in \mathcal{X}^{n \times d}$ with $\mathbf{x} \sim P$ and similarly for \mathbf{Y}_n with $\mathbf{y} \sim Q$. Our goal is to find a set of features \mathcal{S} such that for $i \in \mathcal{S}$, the marginal distributions of the i -th dimension of \mathbf{x} and \mathbf{y} (denoted as $P^{(i)}$ and $Q^{(i)}$) are significantly different, i.e., $P^{(i)} \neq Q^{(i)}$.*

The problem and a solution were initially proposed by Yamada et al. (2019) based on the polyhedral lemma for post-selection inference. Their proposal `mmdInf`, which is referred to as PolyMMD in this paper, first selects a set of k features \mathcal{S}_k using MMD and then tests if each of the selected features' marginal distributions are different. The latter part is performed by conditioning on selecting the whole set \mathcal{S}_k . This form of conditioning can be written equivalently as a set of linear constraints (Yamada et al., 2019, section 3.1) and as a result, it is possible to employ the polyhedral

lemma and obtain a truncated normal as their asymptotic null distribution. However, we can relax the conditioning further. Notice that the goal is to test each feature $i \in \mathcal{S}_k$ separately. Thus, given a significance level α , it is sufficient to require type-I error to be no larger than α , conditioned only on $i \in \mathcal{S}_k$, rather than on the full set \mathcal{S}_k . Following Liu et al. (2018), we call this event the minimal conditioning set. While the selection event of PolyMMD can be written as a single polyhedron, the selection event $i \in \mathcal{S}_k$ is more complicated.

In this section, we propose MultiMMD a more powerful variant of PolyMMD by conditioning on the minimal conditioning set. We show how the statistical test can be performed using multiscale bootstrap. Although for the remainder of the section we focus on the incomplete estimator $\widehat{\text{MMD}}_{Inc}^2$, a similar procedure can be applied to the block estimator $\widehat{\text{MMD}}_{Blo}^2$ and the linear time estimator $\widehat{\text{MMD}}_l^2$.

3.1 Proposal: MultiMMD

From the index set of all features \mathcal{I} , MultiMMD finds a subset of k features, denoted by $\mathcal{S}_k \subseteq \mathcal{I}$, that differentiates samples from P and Q . The k features are selected as the k dimensions with the highest scores measured by an estimator of MMD. More precisely, we have $\mathcal{S}_k = \mathcal{S}_{k-1} \cup \{\arg \max_{i \in \mathcal{I} \setminus \mathcal{S}_{k-1}} \widehat{\text{MMD}}_{Inc}^2(\mathbf{X}_n^{(i)}, \mathbf{Y}_n^{(i)})\}$ where $\mathbf{X}_n^{(i)} = [\mathbf{x}_1^{(i)}, \dots, \mathbf{x}_n^{(i)}]^\top$ and $\mathbf{x}^{(i)}$ is the i -th dimension of the random variable \mathbf{x} (and similarly for $\mathbf{Y}_n^{(i)}$) and $\mathcal{S}_0 = \emptyset$.

The selection procedure mentioned above for MultiMMD is the same as PolyMMD, but the statistical test we perform is different. For each selected feature $i \in \mathcal{S}_k$, the hypothesis test we execute is

$$\begin{aligned} H_{0,i} &: \text{MMD}^2(P^{(i)}, Q^{(i)}) = 0 \mid i \in \mathcal{S}_k \text{ is selected,} \\ H_{1,i} &: \text{MMD}^2(P^{(i)}, Q^{(i)}) > 0 \mid i \in \mathcal{S}_k \text{ is selected.} \end{aligned}$$

In contrast to PolyMMD, when testing $H_{0,i}$ for some $i \in \mathcal{S}_k$, all the other selected variables in $\mathcal{S}_k \setminus \{i\}$ are not considered for conditioning. The justification of why tests that condition on the minimal conditioning set is more powerful can be found in Fithian et al. (2014). The main idea is that the monotonicity of the selective type I error, defined as $\mathbb{P}(\text{Reject } H_{0,i} \leq \alpha \mid H_{0,i}, i \in \mathcal{S}_k)$, suggests that we will lose power when we move from coarse selection variables to finer selection variables (Fithian et al., 2014, Proposition 3).

Multiscale bootstrap's flexibility in representing hypothesis and selective regions with indicator functions makes it a suitable candidate to calculate p -values. It requires us to define a set \mathcal{M}

with each of its members $n' \in \mathcal{M}$ specifying the number of elements to be resampled from \mathbf{X}_n and \mathbf{Y}_n (denoted as $\mathbf{X}_{n'}$ and $\mathbf{Y}_{n'}$). We generate bootstrap replicates of statistic $\hat{\mathbf{T}}_n(\mathbf{X}_{n'}, \mathbf{Y}_{n'}) := \sqrt{l_n} [\widehat{\text{MMD}}_{\text{Inc}}^2(\mathbf{X}_{n'}^{(1)}, \mathbf{Y}_{n'}^{(1)}), \dots, \widehat{\text{MMD}}_{\text{Inc}}^2(\mathbf{X}_{n'}^{(d)}, \mathbf{Y}_{n'}^{(d)})]^\top$ where $l_n = |\mathcal{D}_n|$ is the denominator used for $\widehat{\text{MMD}}_{\text{Inc}}^2(\mathbf{X}_n^{(i)}, \mathbf{Y}_n^{(i)})$. For each $n' \in \mathcal{M}$, the statistic is computed B times and the bootstrap probability $\text{BP}_{\gamma^2}(\cdot)$ of the hypothesis region H (and selective region S) is the average number of the B samples that falls within H (and S). As a result, we require a sampler for $\hat{\mathbf{T}}_n(\mathbf{X}_{n'}, \mathbf{Y}_{n'})$ for all $n' \in \mathcal{M}$ and a function that describes whether the statistic falls within the regions H and S . Finally, two linear regression models are fitted: one for H and one for S denoted as $\varphi_H(\cdot)$ and $\varphi_S(\cdot)$ respectively. Assuming that the boundary surfaces can be represented by a polynomial of degree 3, then the existing theory recommends a linear model (Shimodaira, 2008, Section 5.4). The model's predictor variable is the ratio $\gamma_{n'}^2 := \frac{n}{n'}$ and its response variable is $\gamma_{n'} \bar{\Phi}^{-1}(\text{BP}_{\gamma_{n'}^2}(\cdot)) = \beta_0 + \gamma_{n'}^2 \beta_1$ where β_0 is the signed distance from our statistic to the boundary of the region and β_1 the mean curvature at the boundary.

Algorithm 2 MultiMMD($\mathbf{X}_n, \mathbf{Y}_n, k, \mathcal{M}$): Selective p -values for the null hypothesis $H_{0,i}$: $\text{MMD}^2(P^{(i)}, Q^{(i)}) = 0 \mid i \in \mathcal{S}_k$ is selected.

- 1: $\hat{\mathbf{T}}_n(\mathbf{X}_n, \mathbf{Y}_n), \hat{\Sigma} \leftarrow \text{EstimateParam}(\mathbf{X}_n, \mathbf{Y}_n)$
 - 2: $\mathcal{S}_k \leftarrow$ the indexes of k largest values of $\{\widehat{\text{MMD}}^2(\mathbf{X}_n^{(i)}, \mathbf{Y}_n^{(i)})\}_{i \in \mathcal{I}}$
 - 3: **for** $i \in \mathcal{S}_k$ **do**
 - 4: **for** $n' \in \mathcal{M}$ **do**
 - 5: $\gamma_{n'}^2 \leftarrow \frac{n}{n'}$
 - 6: Sample $\{\mathbf{y}_i^*\}_{i=1}^B \stackrel{i.i.d.}{\sim} \mathcal{N}(\hat{\mathbf{T}}_n(\mathbf{X}_n, \mathbf{Y}_n), \gamma_{n'}^2 \hat{\Sigma})$
 - 7: $\text{BP}_{\gamma_{n'}^2}(S) \leftarrow \sum_{i=1}^B \mathbb{1}_S^{(i)}(\mathbf{y}_i^*)/B$
 - 8: **end for**
 - 9: Fit a linear model $\varphi_S(\gamma^2)$ such that $\varphi_S(\gamma^2) = \gamma \bar{\Phi}^{-1}(\text{BP}_{\gamma^2}(S))$.
 - 10: $\hat{\beta}_0^{(i)} \leftarrow \hat{\sigma}_i^{-1} \sqrt{l_n} \widehat{\text{MMD}}_{\text{Inc}}^2(\mathbf{X}_n^{(i)}, \mathbf{Y}_n^{(i)})$
 - 11: $p_i \leftarrow \bar{\Phi}(\hat{\beta}_0^{(i)})/\bar{\Phi}(\hat{\beta}_0^{(i)} + \varphi_S(0))$
 - 12: **end for**
 - 13: **return** $\{p_i\}_{i=0}^k$ and \mathcal{S}_k
-

We begin by describing how to obtain samples of $\hat{\mathbf{T}}_n(\mathbf{X}_{n'}, \mathbf{Y}_{n'})$ for all $n' \in \mathcal{M}$. Suppose that the bootstrap resamples can be represented using the distribution $\mathcal{N}(\hat{\mathbf{T}}_n(\mathbf{X}_n, \mathbf{Y}_n), \hat{\Sigma})$ where $\hat{\Sigma}$ be the sample covariance of $\widehat{\text{MMD}}_{\text{Inc}}^2$, i.e., $\hat{\Sigma} := \frac{1}{l_n-1} \sum_{(i,j) \in \mathcal{D}_n} [\mathbf{h}(\mathbf{z}_i, \mathbf{z}_j) - \bar{\mathbf{h}}][\mathbf{h}(\mathbf{z}_i, \mathbf{z}_j) - \bar{\mathbf{h}}]^\top$ where (recall that $h(\cdot, \cdot)$ is the U-statistic kernel) $\mathbf{h}(\mathbf{z}_i, \mathbf{z}_j) := [h(\mathbf{z}_i^{(1)}, \mathbf{z}_j^{(1)}), \dots, h(\mathbf{z}_i^{(d)}, \mathbf{z}_j^{(d)})]^\top \in \mathbb{R}^d$ and $\bar{\mathbf{h}} :=$

$\frac{1}{l_n} \sum_{(i,j) \in \mathcal{D}_n} \mathbf{h}(\mathbf{z}_i, \mathbf{z}_j)$. The choice of normal distribution is justified as $\hat{\mathbf{T}}_n(\mathbf{X}_n, \mathbf{Y}_n)$ tends to be normally distributed as $n \rightarrow \infty$ (Yamada et al., 2019, Theorem 5). In order to replicate samples of $\hat{\mathbf{T}}_n(\mathbf{X}_{n'}, \mathbf{Y}_{n'})$, notice that its asymptotic distribution is $\hat{\mathbf{T}}_n(\mathbf{X}_{n'}, \mathbf{Y}_{n'}) \sim \mathcal{N}(\hat{\mathbf{T}}_n(\mathbf{X}_n, \mathbf{Y}_n), \frac{l_n}{l_{n'}} \hat{\Sigma})$. For each n' , instead of resampling n' elements from \mathbf{X}_n and \mathbf{Y}_n for calculating $\hat{\mathbf{T}}_n(\mathbf{X}_{n'}, \mathbf{Y}_{n'})$ B times, we generated B replicates directly from $\mathcal{N}(\hat{\mathbf{T}}_n(\mathbf{X}_n, \mathbf{Y}_n), \frac{l_n}{l_{n'}} \hat{\Sigma})$ which is then used to calculate bootstrap probabilities. The former is an $\mathcal{O}(n'B)$ process while the latter is $\mathcal{O}(B)$. In practice, the B replicates are sampled from $\mathcal{N}(\hat{\mathbf{T}}_n(\mathbf{X}_n, \mathbf{Y}_n), \frac{n}{n'} \hat{\Sigma})$ instead because we let l_n and $l_{n'}$ be the typical choice of $l_n = rn$ and $l_{n'} = rn'$ where r is fixed apriori with $0 < r < \infty$. The choice of r affects the distribution of $\hat{\mathbf{T}}_n(\cdot, \cdot)$ (Yamada et al., 2019, section 4). When r is high, its asymptotic distribution tends towards its complete counterpart (i.e., infinite sums of weighted chi-squared variables). But when r is small, it is normally distributed.

For each n' , the B replicates are used to calculate bootstrap probabilities for both the hypothesis region $\text{BP}_{\frac{n}{n'}}(H)$ and the selective region $\text{BP}_{\frac{n}{n'}}(S)$. Note that the hypothesis $H_{0,i}$ can be written as the region $H = \{y \in \mathbb{R}^d : y^{(i)} \leq 0\}$ which has a flat boundary. This means that the curvature is $\beta_1 = 0$ and multiscale bootstrap is not needed for the hypothesis region H . In fact, the signed distance for testing $i \in \mathcal{S}_k$ is $\hat{\beta}_0^{(i)} = \hat{\sigma}_i^{-1} \sqrt{l_n} \widehat{\text{MMD}}_{\text{Inc}}^2(\mathbf{X}_n^{(i)}, \mathbf{Y}_n^{(i)})$ where $\hat{\sigma}_i^2$ is the i th diagonal element of $\hat{\Sigma}$. We have $\varphi_H(\gamma^2) = \hat{\beta}_0^{(i)}$ as a constant function.

However, it is not as easy for the selective region S which requires the application of multiscale bootstrap. S is represented by an indicator function $\mathbb{1}_S^{(i)}(\mathbf{y}^*) = \begin{cases} 1 & \text{if } i \in \mathcal{S}_k^* \\ 0 & \text{if } i \notin \mathcal{S}_k^* \end{cases}$ where \mathcal{S}_k^* is the selected set of k features where our selection algorithm is applied to \mathbf{y}^* . Let $\gamma_{n'}^2 = \frac{n}{n'}$ then, the bootstrap probability is given by $\text{BP}_{\gamma_{n'}^2}(S) = \sum_{i=1}^B \mathbb{1}_S^{(i)}(\mathbf{y}_i^*)/B$ where $\{\mathbf{y}_i^*\}_{i=1}^B \stackrel{i.i.d.}{\sim} \mathcal{N}(\hat{\mathbf{T}}_n(\mathbf{X}_n, \mathbf{Y}_n), \gamma_{n'}^2 \hat{\Sigma})$. For a given $\mathcal{M} = \{n'\}$, we have $|\mathcal{M}|$ pairs of predictor and response $\{(\gamma_{n'}^2, \gamma_{n'} \bar{\Phi}^{-1}(\text{BP}_{\gamma_{n'}^2}(S)))\}_{n' \in \mathcal{M}}$ that is used to fit a linear model $\varphi_S(\gamma^2)$. We define \mathcal{M} to be the set of numbers equally spaced between $0.5n$ to $2n$ in log space with $|\mathcal{M}| = 10$. The function $\varphi_S(\gamma^2)$ can be used to extrapolate to $\gamma^2 = 0$ to obtain the signed distance from our statistic to the boundary of S . Then, our selective p -value for feature $i \in \mathcal{S}_k$ is given by

$$p_i = \bar{\Phi}(\hat{\beta}_0^{(i)})/\bar{\Phi}(\hat{\beta}_0^{(i)} + \varphi_S(0)).$$

We reject $H_{0,i}$ if $p_i < \alpha$. The algorithm is described in Algorithm 2.

4 SELECTIVE INFERENCE WITH HSIC

We consider the problem studied in Yamada et al. (2018)

Problem 2. *Given n samples from the joint distribution $\{(\mathbf{x}_i, \mathbf{y}_i)\}_{i=1}^n \stackrel{i.i.d.}{\sim} P_{\mathbf{x}\mathbf{y}}$ on the domain $\mathcal{X}^d \times \mathcal{Y}$, our goal is to find a subset \mathcal{S} of features of \mathbf{x} such that for each $i \in \mathcal{S}$ there is statistically significant dependency between the feature $\mathbf{x}^{(i)}$ and response \mathbf{y} .*

The goal is to decide if there is some dependence between the marginal distribution $\mathbf{x}^{(i)}$ and the response \mathbf{y} . Whereas in MultiMMD, it compares if the difference between the marginal distributions of $\mathbf{x}^{(i)}$ and $\mathbf{y}^{(i)}$ is zero, i.e., if $\text{MMD}^2(\mathbf{x}^{(i)}, \mathbf{y}^{(i)}) = 0$. In Yamada et al. (2018), a solution was proposed for Problem 2 using HSIC to measure the dependency between the two between $\mathbf{x}^{(i)}$ and \mathbf{y} which we call ‘‘PolyHSIC’’ (previously called `hsicInf`). It begins with first selecting k features with the highest HSIC scores and then a test is performed for each feature. But the conditioning is not minimal and suffers from a loss in power.

In this section, we propose a new estimator based on the incomplete U-statistic estimator for HSIC and analyse its asymptotic distribution. We then extend PolyHSIC to ‘‘MultiHSIC’’. The new proposal conditions on the minimal conditioning set which is made possible with the multiscale bootstrap. Although our procedure allows the use of the HSIC block estimator (Zhang et al., 2018), we observe that the convergence of the estimator to the target normal distribution highly depends on the block size and in turn the number of blocks, which can be challenging to set correctly. As a result, the false rejection rate of the test can be difficult to control. See Appendix C.2 for a numerical simulation that illustrates this problem, and Zaremba et al. (2013, page 5) for a discussion on a similar issue in the block estimator for MMD.

4.1 Incomplete HSIC

We propose an estimator of HSIC that behaves as desired with the type I error at size α (unlike the block estimator) when used in conjunction with multiscale bootstrap. It based on the incomplete U-statistic estimator (Blom, 1976; Janson, 1984) defined as

$$\widehat{\text{HSIC}}_{\text{Inc}}(\mathbf{z}) = \frac{1}{l} \sum_{(i,j,q,r) \in \mathcal{D}} h(i, j, q, r)$$

where $\mathbf{z} = [z_1, \dots, z_n]$ of n i.i.d. draws from $z := (x, y) \sim P_{\mathbf{x}\mathbf{y}}$, $l = |\mathcal{D}|$, and \mathcal{D} is the *design* of the matrix and for $\widehat{\text{HSIC}}_{\text{Inc}}$ it is constructed randomly by sampling l terms with replacement from \mathbf{i}_4^n .

The asymptotic distribution of $\widehat{\text{HSIC}}_{\text{Blo}}$ is normal in both cases when y and x are independent and dependent (Zhang et al., 2018, section 3.2). As shown in Corollary 1, it also follows that $\widehat{\text{HSIC}}_{\text{Inc}}$ is asymptotically normal regardless of the presence of the dependency between x and y (in Appendix B, we empirically validate this claim).

Corollary 1 (Asymptotic Distribution of $\widehat{\text{HSIC}}_{\text{Inc}}$). *Assume that $\lim_{n,l \rightarrow \infty} n^{-2}l = 0$ and $0 < \lim_{n,l \rightarrow \infty} n^{-1}l = \lambda < \infty$,*

- *If $X \perp Y$, then $l^{\frac{1}{2}}\widehat{\text{HSIC}}_{\text{Inc}}(\mathbf{z}) \xrightarrow{d} \mathcal{N}(0, \sigma^2)$,*
- *If $X \not\perp Y$, then $l^{\frac{1}{2}}(\widehat{\text{HSIC}}_{\text{Inc}}(\mathbf{z}) - \text{HSIC}(P_{\mathbf{x}\mathbf{y}})) \xrightarrow{d} \mathcal{N}(0, \lambda\sigma_u^2 + \sigma^2)$,*

where $\sigma^2 = \text{Var}[h(i, j, q, r)]$ and σ_u^2 is the variance of the complete U-statistic counterpart, see Song et al. (2012, Theorem 5).

A measure of its performance is asymptotic relative efficiency (ARE) (Lee, 2019) of the incomplete estimator with respect its complete counterpart

$$\begin{aligned} \text{ARE} &= \lim_{n \rightarrow \infty} \text{Var}(\widehat{\text{HSIC}}_{\text{Inc}}) / \text{Var}(\widehat{\text{HSIC}}_u) \\ &= \lim_{n \rightarrow \infty} \frac{(\lambda\sigma_u^2 + \sigma^2)/l}{\sigma_u^2/n} = 1 + \frac{\sigma^2}{r\sigma_u^2}, \end{aligned}$$

where a common choice of l is chosen to be rn for some $r > 0$. This means that for large r the incomplete estimator is asymptotically efficient and dependent on the ratio $\frac{\sigma^2}{\sigma_u^2}$ i.e., if we set r to be big, we have $\text{ARE} \approx 1$. A similar analysis can be performed for the incomplete estimator for the MMD which suggests we should take r to be very high but it would violate our assumption that $\lim_{n \rightarrow \infty} n^{-1}l < \infty$ so the estimator will deviate from normal.

4.2 Proposal: MultiHSIC

In this section, we outline MultiHSIC as a more powerful method for variable selection by considering the minimal conditioning set. The algorithm begins by selecting k features with the k largest $\widehat{\text{HSIC}}$ scores and then performing a test after selection. We have

$$\begin{aligned} H_{0,i} &: \text{HSIC}(P_{\mathbf{x}^{(i)}\mathbf{y}}) = 0 \mid i \in \mathcal{S}_k \text{ is selected,} \\ H_{1,i} &: \text{HSIC}(P_{\mathbf{x}^{(i)}\mathbf{y}}) > 0 \mid i \in \mathcal{S}_k \text{ is selected,} \end{aligned}$$

where $P_{\mathbf{x}^{(i)}\mathbf{y}}$ is defined as the joint distribution between $\mathbf{x}^{(i)}$ and \mathbf{y} . Define $\mathbf{Z}^{(i)} = [z_1^{(i)}, \dots, z_n^{(i)}]$ and $\mathbf{z}^{(i)} := (x^{(i)}, y)$. Let $\mathcal{S}_k = \mathcal{S}_{k-1} \cup \{\arg \max_{i \in \mathcal{D} \setminus \mathcal{S}_{k-1}} \widehat{\text{HSIC}}_{\text{Inc}}(\mathbf{Z}^{(i)})\}$ and $\mathcal{S}_0 = \emptyset$.

Related Work: Slim et al. (2019) proposed a class of kernel based statistics for selecting variables which can later be used for hypothesis testing. Their selective inference algorithm includes a sampler for simulating the null hypothesis. While we have focused on feature selection algorithms that condition on selection events to solve Problem 2, there is another branch of selective inference procedures based on the knockoff filter (Barber et al., 2019) which has been extended to the high dimensional setting (Candes et al., 2018). Their framework is similar in the sense that they too first select promising features and then provide selective guarantee on the inference made on the selected variables. However, their guarantees are based on generating convincing knock-offs variables such that the joint distribution is invariant between swaps of the variables and its knockoffs, which can be hard (Lu et al., 2018; Romano et al., 2019). These proposals control false discovery rate (Benjamini and Hochberg, 1995).

5 EXPERIMENTS

In this section, we demonstrate our proposed method for both toy and real world datasets. The performance of our algorithm is measured by true positive rate (TPR) and false positive rate (FPR) which can be thought of as power and type-I error. TPR is defined to be the portion of true selected features that are correctly declared as such and FPR quantifies the portion of selected false features that are declared as incorrectly significant (see definitions in Section A). It is desirable to have high true positive rate and false positive rate to be controlled at α (it is not desirable for this to be below α or above α) since the threshold is chosen to be such that the type-I error is size α . Unless specified otherwise, we use the Gaussian kernel with its bandwidth chosen with the median heuristic.

Our first experiment considers several synthetic problems to evaluate our proposal and verify that our test controls FPR at nominal levels. For MMD, we use the mean shift problem varying both n and d and, for HSIC, we consider the logistic problem. Then, we proceed to using several real world data-sets which have been augmented with artificial and independent features. We consider the original preprocessed features as “true” features which allows us to calculate TPR and FPR. For MMD, we split the data-set into two sets for two different classes and the goal here is to “rediscover” the original features (with a minimal number of artificially added and uninformative features). And for HSIC, the data-set is split into the predictor variables (with some fakes) and the response variable, the goal here is to find the original predictors. For our final experiment, we consider the problem of anomalous dataset detection where d is small




Figure 1: Mean shift experiment as n increases. Results are shown for $\widehat{\text{MMD}}_l^2$ and $\widehat{\text{MMD}}_{Inc}^2$.




Figure 2: Mean shift experiment as d increases. Results are shown for $\widehat{\text{MMD}}_l^2$ and $\widehat{\text{MMD}}_{Inc}^2$.

(and so k is small too). In this scenario, our algorithm only has incremental increase in power. Additional experiments can be found in the Appendix E. Code for reproducing our results is available online: <https://github.com/jenninglim/multiscale-features>.

5.1 Toy Problems

The aim of these synthetic experiment is to evaluate our proposals, MultiMMD and MultiHSIC, against previously proposed methods and empirically verify the theoretical guarantees. The TPR and FPR are averaged over 100 trials, $k = 30$ and $\alpha = 0.05$. We consider the three scenarios.

MMD: Mean Shift with varying n ($d = 50$). We are given n samples from $P = \mathcal{N}(\mathbf{0}, \mathbf{I})$ and $Q = \mathcal{N}(\boldsymbol{\mu}, \mathbf{I})$ where $\boldsymbol{\mu} = [\mathbf{0.5}_{10}, \mathbf{0}_{40}]^\top \in \mathbb{R}^{50}$. For the first ten rows the alternative holds while for the rest the test should not reject the null hypothesis. This problem was studied in Yamada et al. (2019) and the results are shown in Figure 1.

MMD: Mean Shift with varying d ($n = 1000$). The samples are drawn from $P = \mathcal{N}(\mathbf{0}, \mathbf{I})$ and $Q = \mathcal{N}(\boldsymbol{\mu}, \mathbf{I})$ where $\boldsymbol{\mu} = [\mathbf{0.5}_{10}, \mathbf{0}_{d-10}]^\top \in \mathbb{R}^d$. The alternative holds only for the first ten rows. The results




Figure 3: Logistic experiment as n increases. Results are shown for $\widehat{\text{HSIC}}_{Blo}$ and $\widehat{\text{HSIC}}_{Inc}$.

are shown in Figure 2.

HSIC: Logistic problem with varying n ($d = 50$).

We consider the feature selection toy experiment studied in Jordon et al. (2018); Candès et al. (2016). We have $\mathbf{x} = [x_1, \dots, x_n]$ is n i.i.d. draws from 50-dimensional $\mathcal{N}(\mathbf{0}, \mathbf{I})$ and $\mathbf{y} = [y_1, \dots, y_n]$ with $y_j \sim \text{Bernoulli}(\text{Logistic}(\sum_{i=1}^{10} x_j^{(i)}))$ where $\text{Logistic}(x) = \frac{\exp(x)}{1 + \exp(x)}$. Notice that \mathbf{y} is dependent only on the first 10 dimensions of \mathbf{x} and thus it is desirable to only reject the null hypothesis for these first 10 features. For the block estimator, we set the block size to 5; and for the incomplete estimator, we set $r = 1$. The results are shown in Figure 3.

5.2 Benchmarks

We apply MultiMMD and PolyMMD for selecting features that significantly distinguishes two samples. Since TPR and FPR requires the knowledge of true features which is unknown, we regard the original d' number of pre-processed features in the dataset as “true” features and then we augment the dataset with 30 fake features. This problem was studied by Yamada et al. (2019). We apply our proposal to three datasets.

Pulsar dataset ($n = 100$, $d' = 8$) of Lyon et al. (2016) contain samples of pulsar candidates collected during the High Time Resolution Universe Survey. We split the dataset into two sets where one is for pulsars and the other for not pulsars.

Heart dataset ($n = 138$, $d' = 13$) of Janosi et al. (1988) contains samples of patients, their attributes (such as age and sex) and whether they suffer from heart disease. We split the dataset by whether they have heart disease or not.

Wine dataset ($n = 100$, $d' = 12$) of Cortez et al. (2009) contains samples related to red and white variants of the Portuguese “Vinho Verde” wine. The dataset is split into red and white wines.

	PolyMMD		MultiMMD	
Dataset	TPR	FPR	TPR	FPR
Pulsar	0.746	0.063	0.993	0.056
Heart	0.359	0.042	0.588	0.049
Wine	0.567	0.054	0.749	0.057

Table 1: Benchmarking experiment using MMD_{Inc} . The results are averaged over 100 trials ($\alpha = 0.05$).

The results are shown in Table 1. It can be seen that TPR of our proposed method is higher than PolySel for all datasets while both methods corroborate with the theory that the FPR is controlled at α in all scenarios.

5.3 Anomalous Dataset Detection

In this experiment, we are given 6 datasets with one desired reference set and our goal is to eliminate the datasets that deviate too far from the reference. To be specific, our datasets are formed from the smiling subset of the CelebA dataset (Liu et al., 2015), it may also contain synthetic samples generated from the smiling GAN of Jitkrittum et al. (2018). Instead of testing on raw pixels, the datasets are pre-processed and represented by 2048-dimensional features extracted from the Pool3 layer of Inception-v3 (Szegedy et al., 2016). Each dataset contains 1000 samples with $x\%$ being fake images and $1 - x\%$ real images. In this case, since all models are wrong (Box, 1976), the higher percentage of the presence of synthetic samples, the higher the chance of rejection. We apply MultiSel and PolySel with the IMQ kernel (Gorham and Mackey, 2017).

Dataset	1	2	3	4	5	6
% Fakes	0%	10%	20%	30%	40%	50%
MultiMMD	0.03	0.02	0.07	0.06	0.27	0.49
PolyMMD	0.02	0.02	0.05	0.04	0.28	0.45

Table 2: Rejection rate of five datasets for both MultiSel and PolySel. Each dataset has its own percentage of fake features. These results were averaged over 100 trials and we set $k = 4$, $\alpha = 0.05$ and $n = 2000$.

The results are shown in Table 2. The rejection rates of both methods are similar. Dataset 1 has the same distribution as our reference model and so the rejection rate of less than α . As for the other datasets, the rejection rate increases as the percentage of fake increases but the similarity in the performance is expected and can be explained by the small difference in the selection event for PolyMMD and MultiMMD.

Acknowledgements

M.Y. was supported by the JST PRESTO program JPMJPR165A and partly supported by MEXT KAKENHI 16H06299 and the RIKEN engineering network funding. S.M. was supported by MEXT KAKENHI 16H06299.

References

- Barber, R. F., Candès, E. J., et al. (2019). A knockoff filter for high-dimensional selective inference. *The Annals of Statistics*, 47(5):2504–2537.
- Benjamini, Y. and 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.
- Blom, G. (1976). Some properties of incomplete u-statistics. *Biometrika*, 63(3):573–580.
- Box, G. E. (1976). Science and statistics. *Journal of the American Statistical Association*, 71(356):791–799.
- Candès, E., Fan, Y., Janson, L., and Lv, J. (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.
- Candès, E. J., Fan, Y., Janson, L., and Lv, J. (2016). *Panning for gold: Model-free knockoffs for high-dimensional controlled variable selection*. Department of Statistics, Stanford University.
- Cortez, P., Cerdeira, A., Almeida, F., Matos, T., and Reis, J. (2009). Modeling wine preferences by data mining from physicochemical properties. *Decision Support Systems*, 47(4):547–553.
- Efron, B., Halloran, E., and Holmes, S. (1996). Bootstrap confidence levels for phylogenetic trees. *Proceedings of the National Academy of Sciences*, 93(23):13429–13429.
- Efron, B., Tibshirani, R., et al. (1998). The problem of regions. *The Annals of Statistics*, 26(5):1687–1718.
- Fithian, W., Sun, D., and Taylor, J. (2014). Optimal inference after model selection. *arXiv preprint arXiv:1410.2597*.
- Fithian, W., Taylor, J., Tibshirani, R., and Tibshirani, R. (2015). Selective sequential model selection. *arXiv preprint arXiv:1512.02565*.
- Fukumizu, K., Gretton, A., Sun, X., and Schölkopf, B. (2008). Kernel measures of conditional dependence. In *Advances in neural information processing systems*, pages 489–496.
- Gelman, A. and Loken, E. (2013). The garden of forking paths: Why multiple comparisons can be a problem, even when there is no fishing expedition or p-hacking and the research hypothesis was posited ahead of time. *Department of Statistics, Columbia University*.
- Gorham, J. and Mackey, L. (2017). Measuring sample quality with kernels. In *Proceedings of the 34th International Conference on Machine Learning—Volume 70*, pages 1292–1301. JMLR. org.
- Gretton, A., Borgwardt, K. M., Rasch, M. J., Schölkopf, B., and Smola, A. (2012). A kernel two-sample test. *Journal of Machine Learning Research*, 13(Mar):723–773.
- Gretton, A., Bousquet, O., Smola, A., and Schölkopf, B. (2005). Measuring statistical dependence with Hilbert-Schmidt norms. In *International conference on algorithmic learning theory*, pages 63–77. Springer.
- Gretton, A., Fukumizu, K., Teo, C. H., Song, L., Schölkopf, B., and Smola, A. J. (2008). A kernel statistical test of independence. In *Advances in neural information processing systems*, pages 585–592.
- Hoeffding, W. (1992). A class of statistics with asymptotically normal distribution. In *Breakthroughs in Statistics*, pages 308–334. Springer.
- Janosi, A., Steinbrunn, W., Pfisterer, M., and De-trano, R. (1988). Uci machine learning repository-heart disease dataset.
- Janson, S. (1984). The asymptotic distributions of incomplete U-statistics. *Zeitschrift für Wahrscheinlichkeitstheorie und Verwandte Gebiete*, 66(4):495–505.
- Jitkrittum, W., Kanagawa, H., Sangkloy, P., Hays, J., Schölkopf, B., and Gretton, A. (2018). Informative features for model comparison. In *Advances in Neural Information Processing Systems*, pages 808–819.
- Jordon, J., Yoon, J., and van der Schaar, M. (2018). KnockoffGAN: Generating knockoffs for feature selection using generative adversarial networks.
- Lee, A. J. (2019). *U-statistics: Theory and Practice*. Routledge.
- Lee, J. D., Sun, D. L., Sun, Y., Taylor, J. E., et al. (2016). Exact post-selection inference, with application to the lasso. *The Annals of Statistics*, 44(3):907–927.
- Lim, J. N., Yamada, M., Schölkopf, B., and Jitkrittum, W. (2019). Kernel stein tests for multiple model comparison. In *Advances in Neural Information Processing Systems*, pages 2240–2250.

- Liu, K., Markovic, J., and Tibshirani, R. (2018). More powerful post-selection inference, with application to the Lasso. *arXiv preprint arXiv:1801.09037*.
- Liu, Z., Luo, P., Wang, X., and Tang, X. (2015). Deep learning face attributes in the wild. In *Proceedings of International Conference on Computer Vision (ICCV)*.
- Lu, Y., Fan, Y., Lv, J., and Noble, W. S. (2018). Deep-PINK: reproducible feature selection in deep neural networks. In *Advances in Neural Information Processing Systems*, pages 8676–8686.
- Lyon, R. J., Stappers, B., Cooper, S., Brooke, J., and Knowles, J. (2016). Fifty years of pulsar candidate selection: from simple filters to a new principled real-time classification approach. *Monthly Notices of the Royal Astronomical Society*, 459(1):1104–1123.
- Romano, Y., Sesia, M., and Candès, E. (2019). Deep knockoffs. *Journal of the American Statistical Association*, (just-accepted):1–27.
- Serfling, R. J. (2009). *Approximation theorems of mathematical statistics*, volume 162. John Wiley & Sons.
- Shimodaira, H. (2002). An approximately unbiased test of phylogenetic tree selection. *Systematic biology*, 51(3):492–508.
- Shimodaira, H. (2008). Testing regions with nonsmooth boundaries via multiscale bootstrap. *Journal of Statistical Planning and Inference*, 138(5):1227–1241.
- Shimodaira, H. (2014). Higher-order accuracy of multiscale-double bootstrap for testing regions. *Journal of Multivariate Analysis*, 130:208–223.
- Shimodaira, H. et al. (2004). Approximately unbiased tests of regions using multistep-multiscale bootstrap resampling. *The Annals of Statistics*, 32(6):2616–2641.
- Shimodaira, H. and Terada, Y. (2019). Selective inference for testing trees and edges in phylogenetics. *Frontiers in Ecology and Evolution*, 7:174.
- Simmons, J. P., Nelson, L. D., and Simonsohn, U. (2011). False-positive psychology: Undisclosed flexibility in data collection and analysis allows presenting anything as significant. *Psychological science*, 22(11):1359–1366.
- Slim, L., Chatelain, C., Azencott, C.-A., and Vert, J.-P. (2019). kernelpsi: a post-selection inference framework for nonlinear variable selection. In *International Conference on Machine Learning*, pages 5857–5865.
- Smola, A., Gretton, A., Song, L., and Schölkopf, B. (2007). A Hilbert space embedding for distributions. In *International Conference on Algorithmic Learning Theory*, pages 13–31. Springer.
- Song, L., Smola, A., Gretton, A., Bedo, J., and Borgwardt, K. (2012). Feature selection via dependence maximization. *Journal of Machine Learning Research*, 13(May):1393–1434.
- Szegedy, C., Vanhoucke, V., Ioffe, S., Shlens, J., and Wojna, Z. (2016). Rethinking the inception architecture for computer vision. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 2818–2826.
- Terada, Y. and Shimodaira, H. (2017). Selective inference for the problem of regions via multiscale bootstrap. *arXiv preprint arXiv:1711.00949*.
- Terada, Y. and Shimodaira, H. (2019). Selective inference after variable selection via multiscale bootstrap. *arXiv preprint arXiv:1905.10573*.
- Tian, X., Taylor, J., et al. (2018). Selective inference with a randomized response. *The Annals of Statistics*, 46(2):679–710.
- Tibshirani, R. J., Taylor, J., Lockhart, R., and Tibshirani, R. (2016). Exact post-selection inference for sequential regression procedures. *Journal of the American Statistical Association*, 111(514):600–620.
- Yamada, M., Umezu, Y., Fukumizu, K., and Takeuchi, I. (2018). Post selection inference with kernels. In *Proceedings of the Twenty-First International Conference on Artificial Intelligence and Statistics*, volume 84 of *Proceedings of Machine Learning Research*, pages 152–160, Playa Blanca, Lanzarote, Canary Islands. PMLR.
- Yamada, M., Wu, D., Tsai, Y.-H. H., Ohta, H., Salakhutdinov, R., Takeuchi, I., and Fukumizu, K. (2019). Post selection inference with incomplete maximum mean discrepancy estimator. In *International Conference on Learning Representations*.
- Yöntem, M. K., Adem, K., İlhan, T., and Kılıçarslan, S. (2019). Divorce prediction using correlation based feature selection and artificial neural networks. *Neuşehir Hacı Bektaş Veli Üniversitesi SBE Dergisi*, 9(1):259–273.
- Zaremba, W., Gretton, A., and Blaschko, M. (2013). B-test: A non-parametric, low variance kernel two-sample test. In *Advances in neural information processing systems*, pages 755–763.
- Zhang, Q., Filippi, S., Gretton, A., and Sejdinovic, D. (2018). Large-scale kernel methods for independence testing. *Statistics and Computing*, 28(1):113–130.