

Nonparametric Independence Testing for Small Sample Sizes

Aaditya Ramdas*

Dept. of Statistics and Machine Learning Dept.
Carnegie Mellon University
aramdas@cs.cmu.edu

Leila Wehbe*

Machine Learning Dept.
Carnegie Mellon University
lwehbe@cs.cmu.edu

Abstract

This paper deals with the problem of nonparametric independence testing, a fundamental decision-theoretic problem that asks if two arbitrary (possibly multivariate) random variables X, Y are independent or not, a question that comes up in many fields like causality and neuroscience. While quantities like correlation of X, Y only test for (univariate) linear independence, natural alternatives like mutual information of X, Y are hard to estimate due to a serious curse of dimensionality. A recent approach, avoiding both issues, estimates norms of an operator in Reproducing Kernel Hilbert Spaces (RKHSs). Our main contribution is strong empirical evidence that by employing *shrunk* operators when the sample size is small, one can attain an improvement in power at low false positive rates. We analyze the effects of Stein shrinkage on a popular test statistic called HSIC (Hilbert-Schmidt Independence Criterion). Our observations provide insights into two recently proposed shrinkage estimators, SCOSE and FCOSE - we prove that SCOSE is (essentially) the optimal linear shrinkage method for *estimating* the true operator; however, the nonlinearly shrunk FCOSE usually achieves greater improvements in *test power*. This work is important for more powerful nonparametric detection of subtle nonlinear dependencies for small samples.

1 Introduction

The problem of *nonparametric* independence testing deals with ascertaining if two random variables are independent or not, making no parametric assumptions about their underlying distributions. Formally, given n samples (x_i, y_i) for $i \in \{1, \dots, n\}$ where $x_i \in \mathbb{R}^p, y_i \in \mathbb{R}^q$, that are drawn from a joint distribution P_{XY} supported on $\mathcal{X} \times \mathcal{Y} \subseteq \mathbb{R}^{p+q}$, we want to decide between the *null* and *alternate* hypotheses

$$\mathcal{H}_0 : P_{XY} = P_X \times P_Y \text{ vs. } \mathcal{H}_1 : P_{XY} \neq P_X \times P_Y$$

where P_X, P_Y are the marginals of P_{XY} w.r.t. X, Y . A test is a function from the data to $\{0, 1\}$. Tests aim to have high

power (probability of detecting dependence, when it exists) at a prespecified allowable type-I error rate α (probability of detecting dependence when there isn't any).

Independence testing is often a precursor to further analysis. Consider for instance conditional independence testing for inferring causality, say by the PC algorithm [Spirtes *et al.*, 2000], whose first step is (unconditional) independence testing. It is also useful for scientific discovery like in neuroscience, to see if a stimulus X (say an image) is independent of the brain activity Y (say fMRI) in a relevant part of the brain. Since *detecting* nonlinear correlations is much easier than *estimating* a nonparametric regression function (of Y onto X), it can be done at smaller sample sizes, with further samples collected for estimation only if an effect is detected by the hypothesis test. For such situations, correlation only tests for univariate linear independence, while other statistics like mutual information that do characterize multivariate independence are hard to estimate from data, suffering from a serious curse of dimensionality. A recent popular approach for this problem (and a related two-sample testing problem) involve the use of quantities defined in reproducing kernel Hilbert spaces (RKHSs) - see [Gretton *et al.*, 2006; Harchaoui *et al.*, 2007; Gretton *et al.*, 2005b; 2005a].

This paper will concern itself with increasing the statistical power at small samples of a popular kernel statistic called HSIC, by using *shrunk* empirical estimators of the unknown population quantity (introduced below).

1.1 Hilbert Schmidt Independence Criterion

Due to limited space, familiarity with RKHS terminology is assumed - see [Scholkopf and Smola, 2002] for an introduction. Let $k : \mathcal{X} \times \mathcal{X} \rightarrow \mathbb{R}$ and $l : \mathcal{Y} \times \mathcal{Y} \rightarrow \mathbb{R}$ be two positive-definite reproducing kernels that correspond to RKHSs \mathcal{H}_k and \mathcal{H}_l respectively with inner-products $\langle \cdot, \cdot \rangle_k$ and $\langle \cdot, \cdot \rangle_l$. Let k, l arise from (implicit) feature maps $\phi : \mathcal{X} \rightarrow \mathcal{H}_k$ and $\psi : \mathcal{Y} \rightarrow \mathcal{H}_l$. In other words, ϕ, ψ are not functions, but mappings to the Hilbert space. i.e. $\phi(x) \in \mathcal{H}_k, \psi(y) \in \mathcal{H}_l$ respectively. These functions, when evaluated at points in the original spaces, must satisfy $\phi(x)\langle x' \rangle = \langle \phi(x), \phi(x') \rangle_k = k(x, x')$ and $\psi(y)\langle y' \rangle = \langle \psi(y), \psi(y') \rangle_l = l(y, y')$.

The mean embedding of P_X and P_Y are defined as $\mu_X := \mathbb{E}_{x \sim P_X} \phi(x) \in \mathcal{H}_k$ and $\mu_Y := \mathbb{E}_{y \sim P_Y} \psi(y) \in \mathcal{H}_l$ whose empirical estimates are $\hat{\mu}_X := \frac{1}{n} \sum_{i=1}^n \phi(x_i)$ and $\hat{\mu}_Y := \frac{1}{n} \sum_{i=1}^n \psi(y_i)$. Finally, the cross-covariance operator of

*Both authors have equally contributed to this paper.

X, Y is defined as

$$\Sigma_{XY} := \mathbb{E}_{(x,y) \sim P_{XY}} (\phi(x) - \mu_X) \otimes (\psi(y) - \mu_Y)$$

where \otimes is an outer-product. For unfamiliar readers, if we used the linear kernel $k(x, x') = x^T x'$ and $l(y, y') = y^T y'$, then the cross-covariance operator is just the cross-covariance matrix. The plug-in empirical estimator of Σ_{XY} is

$$S_{XY} := \frac{1}{n} \sum_{i=1}^n (\phi(x_i) - \hat{\mu}_X) \otimes (\psi(y_i) - \hat{\mu}_Y)$$

For conciseness, define $\tilde{\phi}(x_i) = \phi(x_i) - \hat{\mu}_X$, $\tilde{\psi}(y_i) = \psi(y_i) - \hat{\mu}_Y$, $\tilde{k}(x, x') = \langle \tilde{\phi}(x), \tilde{\phi}(x') \rangle_k$ and $\tilde{l}(y, y') = \langle \tilde{\psi}(y), \tilde{\psi}(y') \rangle_l$. The test statistic Hilbert-Schmidt Independence Criterion (HSIC) defined in [Gretton *et al.*, 2005a] is the squared Hilbert-Schmidt norm of S_{XY} , and can be calculated using centered kernel matrices \tilde{K}, \tilde{L} , where $\tilde{K}_{ij} = \tilde{k}(x_i, x_j)$, $\tilde{L}_{ij} = \tilde{l}(y_i, y_j)$, as

$$\text{HSIC} := \|S_{XY}\|_{HS}^2 = \frac{1}{n^2} \text{tr}(\tilde{K}\tilde{L}) \quad (1)$$

For unfamiliar readers, if we used the linear kernel, this just corresponds to the Frobenius norm of the cross-covariance matrix. The most important property is: *when the kernels k, l are “characteristic”, then the corresponding population statistic $\|\Sigma_{XY}\|_{HS}^2$ is zero iff X, Y are independent* [Gretton *et al.*, 2005a]. This gives rise to a natural test - calculate $\|S_{XY}\|_{HS}^2$ and reject the null if it is large.

Examples of characteristic kernels include Gaussian $k(x, x') = \exp\left(-\frac{\|x-x'\|_2^2}{\gamma^2}\right)$ and Laplace $k(x, x') = \exp\left(-\frac{\|x-x'\|_1}{\gamma}\right)$, for any bandwidth γ , while the aforementioned linear kernel is not characteristic — the corresponding HSIC tests only linear relationships, and a zero cross-covariance matrix characterizes independence only for multivariate Gaussian distributions. Working with the infinite dimensional operator with characteristic kernels, allows us to identify any general nonlinear dependence (in the limit) between any pair of distributions, not just Gaussians.

1.2 Independence Testing using HSIC

A permutation-based test is described in [Gretton *et al.*, 2005a], and proceeds in the following manner. From the given data, calculate the test statistic $T := \|S_{XY}\|_{HS}^2$. Keeping the order of x_1, \dots, x_n fixed, randomly permute y_1, \dots, y_n a large number of times, and recompute the *permuted* HSIC each time. This destroyed any dependence between x, y simulating a draw from the product of marginals, making the empirical distribution of the permuted HSICs behave like the null distribution of the test statistic (distribution of HSIC when \mathcal{H}_0 is true). For a pre-specified type-1 error α , calculate threshold t_α in the right tail of the null distribution. Reject \mathcal{H}_0 if $T > t_\alpha$. This test was proved to be *consistent* against any fixed alternative, meaning for any fixed type-1 error α , the power goes to 1 as $n \rightarrow \infty$. Empirically, the power can be calculated using simulations by repeating the above permutation test many times for a fixed P_{XY} (for which dependence

holds), and reporting the empirical probability of rejecting the null (detecting the dependence). Note that the power depends on P_{XY} (unknown to the user of the test).

1.3 Shrunk Estimators of S_{XY}

Even though S_{XY} is an unbiased estimator of Σ_{XY} , it typically has high variance at low sample sizes. The idea of Stein shrinkage [Stein, 1956] is to trade-off bias and variance, first introduced in the context of Gaussian mean estimation. This strategy of introducing some bias and decreasing the variance to get different estimators of Σ_{XY} was followed by [Muandet *et al.*, 2014] who define a linear shrinkage estimator of S_{XY} called SCOSE (Simple Covariance Shrinkage Estimator) and a nonlinear shrinkage estimator called FCOSE (Flexible Covariance Shrinkage Estimator). When we refer to shrunk estimators, we implicitly mean SCOSE and FCOSE. We will describe these briefly in Section 2.

1.4 Contributions

Our first contribution is the following :

1. We provide evidence that employing shrunk estimators of Σ_{XY} , instead of S_{XY} , to calculate the aforementioned test statistic, can increase the power of the associated independence test at low false positive rates, when the sample size is small (there is higher variance in estimating infinite-dimensional operators).

Our second contribution is to analyze the effect of shrinkage on the test statistic, to provide some practical insight.

2. The effect of shrinkage on the test-statistic is very similar to soft-thresholding (see Section 4), shrinking very small statistics to zero, and shrinking other values nearly (but not linearly, and nearly (but not) monotonically).

Our last contribution is an insight on the two estimators considered in this paper, SCOSE and FCOSE.

3. We prove that SCOSE is (essentially, up to lower order terms) the optimal/oracle linear shrinkage estimator with respect to quadratic risk (see Section 5). However, we observe that FCOSE typically achieves higher power than SCOSE. This indicates that it may be useful to search for the optimal estimator in a larger class than linearly shrunk estimators, and also that quadratic loss may not be the right loss function for the purposes of test power.

The rest of this paper is organized as follows. Section 2 introduces SCOSE, FCOSE and their corresponding shrunk test statistics. Section 3 presents illuminating experiments that bring out the statistically significant improvement in power over HSIC. Section 4 conducts a deeper investigation into the effect of shrinkage and proves the oracle optimality of SCOSE under quadratic risk.

2 Shrunk Estimators and Test Statistics

Let $\mathcal{HS}(\mathcal{H}_k, \mathcal{H}_l)$ represent the set of Hilbert-Schmidt operators from \mathcal{H}_k to \mathcal{H}_l . We first note that S_{XY} can be written as the solution to the following optimization problem.

$$S_{XY} := \min_{Z \in \mathcal{HS}(\mathcal{H}_k, \mathcal{H}_l)} \frac{1}{n} \sum_{i=1}^n \left\| \tilde{\phi}(x_i) \otimes \tilde{\psi}(y_i) - Z \right\|_{HS}^2$$

Using this idea [Muandet *et al.*, 2014] suggest the following two shrunk/regularized estimators.

From SCOSE to HSIC^S

This is derived in [Muandet *et al.*, 2014] by solving

$$\min_{Z \in \mathcal{HS}(\mathcal{H}_k, \mathcal{H}_l)} \frac{1}{n} \sum_{i=1}^n \left\| \tilde{\phi}(x_i) \otimes \tilde{\psi}(y_i) - Z \right\|_{HS}^2 + \lambda \|Z\|_{HS}^2$$

and the optimal solution (called SCOSE) is

$$S_{XY}^S := \left(1 - \frac{\lambda}{1 + \lambda}\right) S_{XY}$$

where λ (and hence the shrinkage intensity) is estimated by leave-one-out cross-validation (LOOCV), in closed form as

$$\begin{aligned} \rho^S &:= \left(\frac{\lambda^{CV}}{1 + \lambda^{CV}} \right) \\ &= \frac{\left[\frac{1}{n} \sum_{i=1}^n \tilde{K}_{ii} \tilde{L}_{ii} - \frac{1}{n^2} \sum_{i,j=1}^n \tilde{K}_{ij} \tilde{L}_{ij} \right]}{(n-2) \frac{1}{n^2} \sum_{i,j=1}^n \tilde{K}_{ij} \tilde{L}_{ij} + \frac{1}{n^2} \sum_{i=1}^n \tilde{K}_{ii} \tilde{L}_{ii}} \end{aligned}$$

Observing the expression for λ^{CV} in [Muandet *et al.*, 2014], the denominator can be negative (for example, with the Gaussian kernel for small bandwidths, resulting in a kernel matrix close to the identity). This can cause λ^{CV} to be negative, and ρ^S to be (unintentionally) outside the range $[0, 1]$. Though not discussed in [Muandet *et al.*, 2014], we shall follow the convention that when $\rho^S < 0$, we shall use $\rho^S = 0$ and if $\rho^S > 1$, we use $\rho^S = 1$. Indeed, one can show that $\left(1 - \frac{\lambda}{1+\lambda}\right)_+ S_{XY}$ dominates $\left(1 - \frac{\lambda}{1+\lambda}\right) S_{XY}$ where $(x)_+ = \max\{x, 0\}$. In Section 4, we prove that S_{XY}^S is (essentially) the optimal/oracle linear shrinkage estimator with respect to quadratic risk.

We can now calculate the corresponding shrunk statistic $\text{HSIC}^S = \|S_{XY}^S\|_{HS}^2 =$

$$\left(1 - \frac{\frac{1}{n} \sum_{i=1}^n \tilde{K}_{ii} \tilde{L}_{ii} - \text{HSIC}}{(n-2)\text{HSIC} + \frac{1}{n} \sum_{i=1}^n \tilde{K}_{ii} \tilde{L}_{ii}}\right)_+^2 \text{HSIC} \quad (2)$$

While the above expression looks daunting, one thing to note is that the amount that HSIC is shrunk (i.e. the multiplicative factor) depends on the value of HSIC. As we shall see in section 4, small HSIC values get shrunk to zero, but as can be seen above, the shrinkage of HSIC is non-monotonic.

From FCOSE to HSIC^F

The Flexible Covariance Shrinkage Estimator is derived by relying on the Representer theorem, see [Scholkopf and Smola, 2002], to instead minimize

$$\frac{1}{n} \sum_{i=1}^n \left\| \tilde{\phi}(x_i) \otimes \tilde{\psi}(y_i) - \sum_{i=1}^n \frac{\beta_i}{n} \tilde{\phi}(x_i) \otimes \tilde{\psi}(y_i) \right\|_{HS}^2 + \lambda \|\beta\|_2^2$$

over all $\beta \in \mathbb{R}^n$, and the optimal solution (called FCOSE) is

$$S_{XY}^F := \sum_{i=1}^n \frac{\beta_i^\lambda}{n} \tilde{\phi}(x_i) \otimes \tilde{\psi}(y_i)$$

$$\text{where } \beta^\lambda = (\tilde{K} \circ \tilde{L} + \lambda I)^{-1} \tilde{K} \circ \tilde{L}$$

where \circ denotes elementwise (Hadamard) product, $\mathbf{1}$ is the vector $[1, 1, \dots, 1]^T$, and as before the best λ is determined by LOOCV. The procedure to evaluate the optimal λ efficiently is described by [Muandet *et al.*, 2014] - a single eigenvalue decomposition of $\tilde{K} \circ \tilde{L}$ costing $O(n^3)$ can be done, following which evaluating LOOCV is only $O(n^2)$ per λ , see [Muandet *et al.*, 2014], section 3.1 for more details. As before, after picking the λ by LOOCV, we can derive the corresponding shrunk test statistic as

$$\begin{aligned} \text{HSIC}^F &= \|S_{XY}^S\|_{HS}^2 \\ &= \frac{1}{n^2} \text{tr}(M(M + \lambda I)^{-1} M(M + \lambda I)^{-1} M) \end{aligned}$$

where $M = \tilde{K} \circ \tilde{L}$. Note here that the shrinkage is not linear, and the effect on HSIC cannot be seen immediately. Similar to SCOSE, we shall see in section 4, small HSIC values get shrunk to zero (LOOCV chooses a large λ).

3 Linear Shrinkage and Quadratic Risk

In this section, we prove that SCOSE is (essentially) optimal within a particular class of estimators. Such ‘oracle’ arguments also exist elsewhere in the literature, like [Ledoit and Wolf, 2004], so we provide only a brief proof outline.

Proposition 1. *The oracle (with respect to quadratic risk) linear shrinkage estimator and intensity is defined as*

$$S^*, \rho^* := \arg \min_{Z \in \mathcal{HS}, Z = (1-\rho)S_{XY}, 0 \leq \rho \leq 1} \|Z - \Sigma_{XY}\|_{HS}^2$$

and is given by $S^* := (1 - \rho^*)S_{XY}$ where

$$\rho^* := \frac{\mathbb{E}\|S_{XY} - \Sigma_{XY}\|_{HS}^2}{\mathbb{E}\|S_{XY}\|^2}$$

Proof. Define $\alpha^2 = \|\Sigma_{XY}\|_{HS}^2$, $\beta^2 = \mathbb{E}\|S_{XY} - \Sigma_{XY}\|_{HS}^2$, $\delta^2 = \mathbb{E}\|S_{XY}\|^2$. Since $\mathbb{E}[S_{XY}] = \Sigma_{XY}$, it is easy to verify that $\alpha^2 + \beta^2 = \delta^2$. Substituting and expanding the objective, we get:

$$\begin{aligned} \mathbb{E}\|Z - \Sigma_{XY}\|_{HS}^2 &= \mathbb{E}\|-\rho S_{XY} + (S_{XY} - \Sigma_{XY})\|_{HS}^2 \\ &= \rho^2 \delta^2 + \beta^2 - 2\rho(\delta^2 - \alpha^2) \\ &= \rho^2 \alpha^2 + (1 - \rho)^2 \beta^2 \end{aligned}$$

Differentiating and equating to zero, gives $\rho^* = \frac{\beta^2}{\delta^2}$. \square

This ρ^* appears in terms of quantities that depend on the unknown underlying distribution (hence the term *oracle* estimator). We use plugin estimates b, d for β, δ . Let $d^2 = \|S_{XY}\|_{HS}^2 = \frac{1}{n^2} \sum_{i,j=1}^n \tilde{K}_{ij} \tilde{L}_{ij} = \text{HSIC}$. Since β^2 is the variance of S_{XY} , let b^2 be the sample variance of S_{XY} , i.e. $b^2 = \frac{1}{n} \frac{1}{n} \sum_{k=1}^n \|\tilde{\phi}(x_i) \otimes \tilde{\psi}(y_i) - S_{XY}\|^2 = \frac{1}{n} \left[\frac{1}{n} \sum_{i=1}^n \tilde{K}_{ii} \tilde{L}_{ii} - \frac{1}{n^2} \sum_{i,j=1}^n \tilde{K}_{ij} \tilde{L}_{ij} \right]$. Plugging these into S^* and simplifying, we see that $\text{HSIC}^* := \|S^*\|_{HS}^2$ is

$$\text{HSIC}^* = \left(1 - \frac{\frac{1}{n} \sum_{i=1}^n \tilde{K}_{ii} \tilde{L}_{ii} - \text{HSIC}}{n \text{HSIC}}\right)^2 \text{HSIC} \quad (3)$$

Comparing Eq.(3) with Eq.(2) shows that SCOSE is essentially S^* , up to a factor in the denominator which is of the same order as the bias of the HSIC empirical estimator¹ (see Theorem 1 in [Gretton *et al.*, 2005a]). In other words, SCOSE just corresponds to using a slightly different estimator for δ^2 than the simple plugin d^2 , which varies on the same order as the bias $\delta^2 - \mathbb{E}d^2$. Hence SCOSE, as estimated via regularization and LOOCV, is (essentially) the optimal linear shrinkage estimator under quadratic risk.

To the best of our knowledge, this is the first such characterization of *optimality of an estimator achieved through leave-one-out cross-validation*. We are only able to prove this because one can explicitly calculate both the oracle linear shrinkage intensity ρ^* as well as the optimal λ^{CV} (as mentioned in Section 2). This raises a natural open question — can we find other situations where the LOOCV estimator is optimal with respect to some risk measure? (perhaps when explicit calculations are not possible, like ridge regression).

4 Experiments

In this section, we run three kinds of experiments: a) to verify that SCOSE has better quadratic risk than FCOSE and original sample estimator, b) detailed synthetic experiments to verify that shrinkage does improve power, across interesting regimes of $\alpha = \{0.01, 0.05, 0.1\}$, and c) real data obtained from MNIST, to show that we shrinkage detect dependence at much lower samples than the original data size.

4.1 Quadratic Risk

Figure 1 shows that SCOSE is indeed much better than both S_{XY} and FCOSE with respect to quadratic risk. Here, we calculate $\mathbb{E}\|Z - \Sigma_{XY}\|_{HS}^2$ for the distribution given in dataset (A) for $Z \in \{S_{XY}, S_{XY}^S, S_{XY}^F\}$. The expectation is calculated by repeating the experiment 1000 times. Each time Z is calculated according to $N \in \{20, 50, 100\}$ samples and Σ_{XY} is approximated by the empirical cross-covariance matrix on 5,000 samples. The four panels use four different kernels which are linear, polynomial, Laplace and Gaussian from top to bottom. The shrunk estimators are always better than the unshrunk, with a larger difference between SCOSE and FCOSE for finite-dimensional feature spaces (top two). In infinite-dimensional feature spaces (bottom two), SCOSE and FCOSE are much better than the unshrunk estimator but very similar to each other. The differences between all estimators decreases with increasing n , since the sample cross-covariance operator itself becomes very accurate.

4.2 Synthetic Data

We perform synthetic experiments in a wide variety of settings to demonstrate that the shrunk test statistics achieve higher power than HSIC in a variety of settings. We follow the schema provided in the introduction for independence testing and calculating power. We only consider difficult distributions with nonlinear dependence between X, Y , on which linear methods like correlation are shown to fail to

¹HSIC and $\text{HSIC} - 2\text{HSIC}/n - C/n^2$ both converge to population HSIC at same rate determined by the dominant term (HSIC).




Figure 1: All panels show quadratic risk $\mathbb{E}\|X - \Sigma_{XY}\|_{HS}^2$ for $X \in \{S_{XY}, S_{XY}^S, S_{XY}^F\}$. Dataset (A) was used in all four panels, but the kernels were varied - from top to bottom is the linear, quadratic, Gaussian and Laplace kernel.

detect dependence (some of them were used in previous papers on independence testing like [Gretton *et al.*, 2007] and [Chwialkowski and Gretton, 2014]).

For all experiments, $\alpha \in \{0.01, 0.05, 0.1\}$ is chosen as the type-1 error (for choosing the threshold level of the null distribution's right tail). For every setting of parameters of each experiment, power is calculated as the percentage of rejection over 200 repetitions (independent trials), with 2000 permutations per repetition (permutation testing to find the null distribution threshold at level α). We use the Gaussian kernel where the bandwidth is chosen by the common median heuristic [Scholkopf and Smola, 2002].

Table 1 is a representative sample from what we saw on other examples - either large, small or no improvement in power was seen but almost never a worsening of power. The improvements in power may not always be huge, but they are statistically significant - it is difficult to detect such non-linear dependencies at low sample sizes, so *any* increase in power can be important in scientific applications.

Remark. A more appropriate way than using error bars to assess significance is by the Wilcoxon rank sum test, omitted for lack of space, though it yields more favorable results.

4.3 Real Data

We use two real datasets - the first is a good example where shrinkage helps a lot, but in the second it does not help (we show it on purpose). Like the synthetic datasets, for most real datasets it either helps or does not hurt (being very rarely worse; see remark in the discussion section).

The first is the Eckerle dataset [Eckerle, 1979] from the NIST Statistical Reference Datasets (NIST StRD) for Non-linear Regression, data from a NIST study of circular interference transmittance ($n=35$, Y is transmittance, X is wavelength). A plot of the data in Figure 2 reveals a nonlinear relationship between X, Y (though the correlation is 0.035 with p-value 0.84). We subsample the data to see how often we can detect a relationship at 10%, 20%, 30% of the original data size, when the false positive level is always controlled at 0.05. The second is the Aircraft dataset [Bowman and Azzalini, 2014] ($n=709$, X is $\log(\text{speed})$, Y is $\log(\text{span})$). Once again, correlation is low, with a p-value of over 0.8, and we subsample the data to 5%, 10%, 20% of the original data size.

















	$\alpha = 0.01$					$\alpha = 0.05$					$\alpha = 0.10$				
	HSIC	HSIC ^S	HSIC ^F			HSIC	HSIC ^S	HSIC ^F			HSIC	HSIC ^S	HSIC ^F		
	0.22 ±0.03	0.21 ±0.03	0.34 ±0.03		✓	0.52 ±0.04	0.52 ±0.04	0.71 ±0.03		✓	0.73 ±0.03	0.72 ±0.03	0.90 ±0.02		✓
	0.41 ±0.03	0.41 ±0.03	0.48 ±0.04		✓	0.68 ±0.03	0.68 ±0.03	0.88 ±0.02		✓	0.85 ±0.03	0.85 ±0.02	0.99 ±0.01		✓
	0.41 ±0.03	0.40 ±0.03	0.52 ±0.04		✓	0.74 ±0.03	0.74 ±0.03	0.94 ±0.02		✓	0.94 ±0.02	0.94 ±0.02	0.99 ±0.01		✓
	0.52 ±0.04	0.52 ±0.04	0.66 ±0.03		✓	0.91 ±0.02	0.91 ±0.02	0.89 ±0.02			0.99 ±0.01	0.99 ±0.01	0.96 ±0.01		✗
	0.04 ±0.01	0.04 ±0.01	0.04 ±0.01			0.12 ±0.02	0.12 ±0.02	0.14 ±0.02			0.23 ±0.03	0.23 ±0.03	0.24 ±0.03		
	0.10 ±0.02	0.10 ±0.02	0.12 ±0.02			0.31 ±0.03	0.31 ±0.03	0.40 ±0.03		✓	0.47 ±0.04	0.47 ±0.04	0.58 ±0.03		✓
	0.33 ±0.03	0.33 ±0.03	0.46 ±0.04		✓	0.77 ±0.03	0.77 ±0.03	0.91 ±0.02		✓	0.95 ±0.01	0.96 ±0.01	0.99 ±0.01		✓
	0.93 ±0.02	0.93 ±0.02	0.96 ±0.01		✓	1.00 ±0.00	1.00 ±0.00	1.00 ±0.00			1.00 ±0.00	1.00 ±0.00	1.00 ±0.00		
	0.07 ±0.02	0.07 ±0.02	0.09 ±0.02			0.24 ±0.03	0.26 ±0.03	0.32 ±0.03		✓	0.44 ±0.04	0.47 ±0.04	0.48 ±0.04		
	0.06 ±0.02	0.07 ±0.02	0.09 ±0.02			0.26 ±0.03	0.28 ±0.03	0.32 ±0.03			0.45 ±0.04	0.47 ±0.04	0.48 ±0.04		
	0.10 ±0.02	0.12 ±0.02	0.14 ±0.02			0.34 ±0.03	0.34 ±0.03	0.39 ±0.03			0.51 ±0.04	0.52 ±0.04	0.53 ±0.04		
	0.07 ±0.02	0.07 ±0.02	0.10 ±0.02		✓	0.30 ±0.03	0.33 ±0.03	0.35 ±0.03			0.53 ±0.04	0.54 ±0.04	0.57 ±0.04		
	0.04 ±0.01	0.05 ±0.02	0.04 ±0.01			0.18 ±0.03	0.27 ±0.03	0.24 ±0.03	✓	✓	0.34 ±0.03	0.45 ±0.04	0.44 ±0.04	✓	✓
	0.16 ±0.03	0.20 ±0.03	0.20 ±0.03			0.45 ±0.04	0.58 ±0.03	0.58 ±0.03	✓	✓	0.67 ±0.03	0.73 ±0.03	0.73 ±0.03	✓	✓
	0.34 ±0.03	0.43 ±0.04	0.43 ±0.04	✓	✓	0.71 ±0.03	0.80 ±0.03	0.79 ±0.03	✓	✓	0.85 ±0.03	0.90 ±0.02	0.89 ±0.02	✓	
	0.63 ±0.03	0.72 ±0.03	0.73 ±0.03	✓	✓	0.91 ±0.02	0.92 ±0.02	0.92 ±0.02			0.95 ±0.01	0.96 ±0.01	0.96 ±0.01		

Table 1: The first column shows scatterplots of X vs Y (all having dependence between X, Y). There are 3 sets of 5 columns each - for $\alpha = 0.01, 0.05, 0.1$ (controlled by running 2000 permutations). In each set, the first three columns show the power of HSIC, HSIC^S, HSIC^F (with standard deviation over 200 repetitions below). The fourth column shows when HSIC^S is significantly better than HSIC, and the fifth column when HSIC^F has significantly higher power than HSIC. A blank means the powers are not significantly better or worse. In the first dataset (A) (top 4) we show how the power varies with increasing n (becomes easier). In the second dataset (B) (second 4) we show how the power varies with rotation (goes from near-independence to clear dependence). In the third dataset (C) (third 4), we demonstrate a case where shrinkage does *not* help much, which is a circle with a hole. In the last dataset (D) (last 4), we demonstrate a case where HSIC^S does as well as HSIC^F. We tried many more datasets, these are a few representative samples.

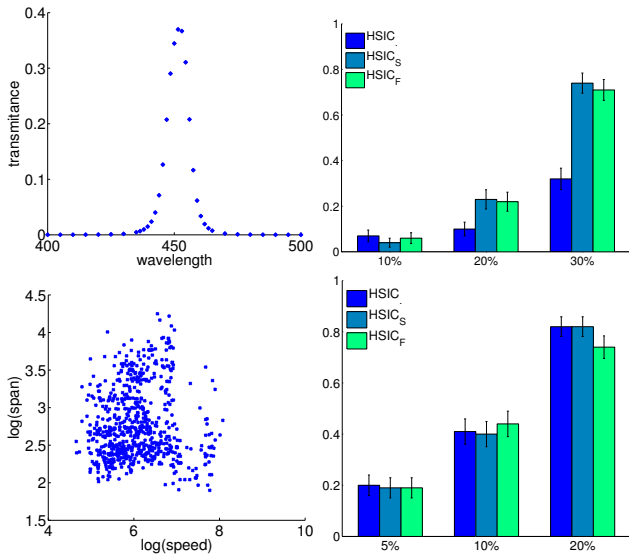


Figure 2: Top Row: The left figure shows a plot of wavelength against transmittance. The right figure shows the power of HSIC , HSIC^S , HSIC^F when the data are subsampled to 10%, 20%, 30% (error bars over 100 repetitions). Bottom Row: The left figure shows a plot of $\log(\text{wingspan})$ vs $\log(\text{airspeed})$. The right figure shows the power of HSIC , HSIC^S , HSIC^F when the data are subsampled to 5%, 10%, 20% (error bars over 100 repetitions).

5 Discussion

Why might shrinkage improve power? Let us examine the net effect of using shrunk estimators on the value of HSIC , i.e. let us compare HSIC^S and HSIC^F to HSIC by computing these over all the repetitions of the permutation testing procedure described in the introduction. In Fig. 3, both estimators are visually similar in transforming the actual test statistic. Perhaps the more interesting phenomenon is that Fig. 3 is reminiscent of the graph of a soft-thresholding operator $ST_t(x) = \max\{0, x - t\}$. Intuitively, if the unshrunk HSIC value is small, the shrinkage methods deem it to be “noise” and it is shrunk to zero. Looking at the X-axis scaling of the top and bottom row, the size of the region that gets shrunk to zero decreases with n - as expected, shrinkage has less effect when S_{XY} has low variance). The shrinkage being non-monotone (more so for $n = 20$ than $n = 50$ in Figure 3) is key to achieving an improvement in power.

Using the intuition from the above figure, we can finally piece together why shrinkage may yield benefits. A rejection of \mathcal{H}_0 occurs when the test statistic stands out in the right tail of its null distribution. Typically, when the alternative is true (this is when rejecting the null improves power) the unshrunk test statistics calculated from the permuted samples is smaller than the unshrunk HSIC calculated on the original sample. However, the effect of shrinking the small statistics towards zero, and setting the smallest ones to zero, is that the unpermuted test statistic under the alternative distribution stands out more in the right tail of the null.




Figure 3: The top row corresponds to $n = 20$, and the bottom row has $n = 50$. The left plots compare HSIC^S to HSIC , and the right plots compare HSIC^F to HSIC . Each cross mark corresponds to the shrunk and unshrunk HSIC calculated during a single permutation of a permutation test.

In other words, relative to the unshrunk null distribution and the unshrunk test statistic, the tail of the null distribution is shrunk more towards zero than the unpermuted test statistic, causing the latter to have a higher quantile in the right tail of the former (relative to the quantile before shrinkage). Let us verify this experimentally. In Fig.4 we plot for each of the datasets T to the 95th percentile of the permuted statistics, for $T \in \{\text{HSIC}, \text{HSIC}^S, \text{HSIC}^F\}$. Recall that for dataset (C), we didn’t see much of an improvement in power, but for (A),(B),(D) it is clear from Fig. 4 that the unpermuted statistic is shrunk less than its null distribution’s 95th quantile.

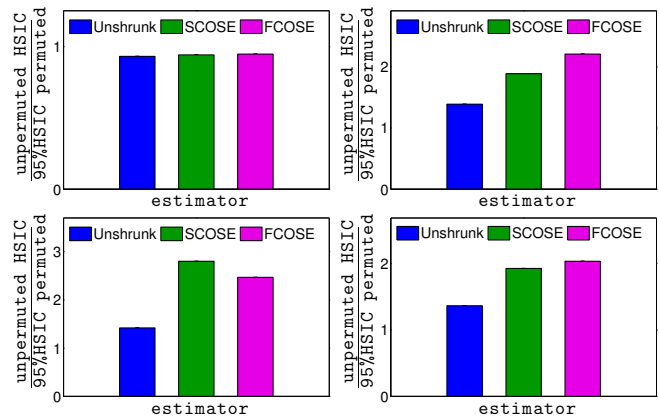


Figure 4: All panels show the ratio of the unpermuted HSIC to the 95th percentile of the null distribution based on HSIC s calculated from the permuted data. (see Table 1) The top row has datasets (C) with radius 2.2, (B) with angle $3 \times \pi/32$, and the bottom row has (D) with $N = 25$, (A) with $N = 40$. These observations were qualitatively the same in all other synthetic data parameter settings, and also for other percentiles than 95th, and since the figures look identical in spirit, they were omitted due to lack of space.

Remark. In our experiments, real and synthetic, shrinkage usually improves (and almost never worsens) power in false-positive regimes that we usually care about. Will shrinkage *always* improve power? Possibly not. Even though shrunk the shrunk S_{XY} dominates S_{XY} for estimation error, it may not be the case that shrunk HSIC always dominates unshrunk HSIC for test power (i.e. the latter may not be *inadmissible*). However, just as no single classifier always outperforms another, it is still beneficial to add techniques like shrinkage, that seem to consistently yield benefits in practice, to the practitioner’s array of tools.

6 Conclusion

We presented evidence for an important phenomenon - using biased but lower variance shrunk estimators of cross-covariance operators can often significantly improve test power of HSIC at small sample sizes. This observation (that shrinkage can improve power) has rarely been made in the statistics and machine learning testing literature. We think the reason is that most test statistics for independence testing cannot be immediately expressed as the norm of an empirical operator, making it less obvious *how* to apply shrinkage to improve their power at low sample sizes.

We also showed the optimality (among linear shrinkage estimators) of SCOSE, but observe that the nonlinear shrinkage of FCOSE usually yields higher power. To the best of our knowledge, there seems to be no current literature showing that the choice made by leave-one-out cross-validation (SCOSE) explicitly leads to an estimator that is “optimal” in some sense (among linear shrinkage estimators). This may be because it is often not possible to explicitly calculate the form of the LOOCV estimator, nor the explicit form of the best linear shrinkage estimator, as can both be done in this simple setting.

Since even the best possible linear shrinkage estimator (as represented by SCOSE) is usually worse than FCOSE, this result indicates that in order to improve upon FCOSE, it will be necessary to further study the class of non-linear shrinkage estimators for our infinite dimensional operators, as done for finite dimensional covariance matrices in [Ledoit and Wolf, 2011] and other papers by the same authors.

We ended with a brief investigation into the effect of shrinkage on HSIC and why shrinkage may intuitively improve power. We think that our work will be important for more powerful nonparametric detection of subtle nonlinear dependencies at low sample sizes, a common problem in scientific applications.

Acknowledgments

We would like to thank Arthur Gretton and Jessica Chemali for useful feedback on an earlier draft of the paper, Larry Wasserman for pointing out some useful references, and Krikamol Muandet for sharing his code. AR was supported in part by ONR MURI grant N000140911052.

References

[Bowman and Azzalini, 2014] A. W. Bowman and A. Azzalini. *R package sm: nonparametric smoothing methods*

(version 2.2-5.4). University of Glasgow, UK and Università di Padova, Italia, 2014.

- [Chwialkowski and Gretton, 2014] Kacper Chwialkowski and Arthur Gretton. A kernel independence test for random processes. In *Proceedings of The 31st International Conference on Machine Learning*, pages 1422–1430, 2014.
- [Eckerle, 1979] K Eckerle. Circular interference transmittance study. *National Institute of Standards and Technology (NIST), US Department of Commerce, USA*, 1979.
- [Gretton *et al.*, 2005a] Arthur Gretton, Olivier Bousquet, Alex Smola, and Bernhard Schölkopf. Measuring statistical dependence with Hilbert-Schmidt norms. In *Algorithmic Learning Theory*, pages 63–77. Springer, 2005.
- [Gretton *et al.*, 2005b] Arthur Gretton, Ralf Herbrich, Alexander Smola, Olivier Bousquet, and Bernhard Schölkopf. Kernel methods for measuring independence. *The Journal of Machine Learning Research*, 6:2075–2129, 2005.
- [Gretton *et al.*, 2006] Arthur Gretton, Karsten M Borgwardt, Malte Rasch, Bernhard Schölkopf, and Alex J Smola. A kernel method for the two-sample-problem. In *Neural Information Processing Systems*, pages 513–520, 2006.
- [Gretton *et al.*, 2007] Arthur Gretton, Kenji Fukumizu, Choon Hui Teo, Le Song, Bernhard Schölkopf, and Alex J Smola. A kernel statistical test of independence. *Neural Information Processing Systems*, 2007.
- [Harchaoui *et al.*, 2007] Zaïd Harchaoui, Francis Bach, and Eric Moulines. Testing for homogeneity with kernel fisher discriminant analysis. *Arxiv preprint 0804.1026*, 2007.
- [Ledoit and Wolf, 2004] Olivier Ledoit and Michael Wolf. A well-conditioned estimator for large-dimensional covariance matrices. *Journal of Multivariate Analysis*, 88(2):365–411, 2004.
- [Ledoit and Wolf, 2011] Olivier Ledoit and Michael Wolf. Nonlinear shrinkage estimation of large-dimensional covariance matrices. *Institute for Empirical Research in Economics University of Zurich Working Paper*, (515), 2011.
- [Muandet *et al.*, 2014] Krikamol Muandet, Kenji Fukumizu, Bharath Sriperumbudur, Arthur Gretton, and Bernhard Schölkopf. Kernel mean estimation and stein effect. In *Proceedings of The 31st International Conference on Machine Learning*, pages 10–18, 2014.
- [Scholkopf and Smola, 2002] Bernhard Scholkopf and Alex Smola. *Learning with kernels*. MIT press Cambridge, 2002.
- [Spirtes *et al.*, 2000] Peter Spirtes, Clark N Glymour, and Richard Scheines. *Causation, prediction, and search*, volume 81. MIT press, 2000.
- [Stein, 1956] Charles Stein. Inadmissibility of the usual estimator for the mean of a multivariate normal distribution. *Proceedings of the Third Berkeley symposium on mathematical statistics and probability*, 1(399):197–206, 1956.