

Approximate Kernel-based Conditional Independence Tests for Fast Non-Parametric Causal Discovery

Eric V. Strobl, Shyam Visweswaran

*Dept. of Biomedical Informatics
University of Pittsburgh*

Kun Zhang

*Dept. of Philosophy
Carnegie Mellon University*

Abstract

Constraint-based causal discovery (CCD) algorithms require fast and accurate conditional independence (CI) testing. The Kernel Conditional Independence Test (KCIT) is currently one of the most popular CI tests in the non-parametric setting, but many investigators cannot use KCIT with large datasets because the test scales cubically with sample size. We therefore devise two relaxations called the Randomized Conditional Independence Test (RCIT) and the Randomized conditional Correlation Test (RCoT) which both approximate KCIT by utilizing random Fourier features. In practice, both of the proposed tests scale linearly with sample size and return accurate p-values much faster than KCIT in the large sample size context. CCD algorithms run with RCIT or RCoT also return graphs at least as accurate as the same algorithms run with KCIT but with large reductions in run time¹.

Keywords: Conditional Independence Test, Non-parametric, Causal Discovery

1. The Problem

Constraint-based causal discovery (CCD) algorithms such as PC and FCI infer causal relations from observational data by combining the results of many conditional independence (CI) tests (Spirtes et al., 2000). In practice, a CCD algorithm can easily request p-values from thousands of CI tests even with a sparse underlying graph. Developing fast and accurate CI tests is therefore critical for maximizing the usability of CCD algorithms across a wide variety of datasets.

¹R implementation at github.com/ericstrobl/RCIT. We recommend that users install Microsoft R Open for fast matrix computations.

Investigators have developed many fast parametric methods for testing CI. For example, we can use partial correlation to test for CI under the assumption of Gaussian variables (Fisher, 1915, 1921). We can also consider testing for unconditional independence $X \perp\!\!\!\perp Y|Z = z, \forall z$ when Z is discrete and $\mathbb{P}(Z = z) > 0$. The chi-squared test for instance utilizes this strategy when both X and Y are also discrete (Pearson, 1900). Another permutation-based test generalizes the same strategy even when X and Y are not necessarily discrete (Tsamardinos and Borboudakis, 2010).

Testing for CI in the non-parametric setting generally demands a more sophisticated approach. One strategy involves discretizing continuous conditioning variables Z as \tilde{Z} in some optimal fashion and assessing unconditional independence $\forall \tilde{Z} = \tilde{z}$ (Margaritis, 2005, Huang, 2010). Discretization however suffers severely from the curse of dimensionality because consistency arguments demand smaller bins with increasing sample size, but the number of cells in the associated contingency table increases exponentially with the conditioning set size. A second method involves measuring the distance between estimates of the conditional densities $f(X|Y, Z)$ and $f(X|Z)$, or their associated characteristic functions, by observing that $f(X|Y, Z) = f(X|Z)$ when $X \perp\!\!\!\perp Y|Z$ (Su and White, 2007, 2008). However, the power of these tests also deteriorates quickly with increases in the dimensionality of Z .

Several investigators have since proposed reproducing kernel-based CI tests in order to tame the curse of dimensionality. Indeed, kernel-based methods in general are known for their strong empirical performance in the high dimensional setting. The Kernel Conditional Independence Test (KCIT) for example assesses CI by capitalizing on a characterization of CI in reproducing kernel Hilbert spaces (RKHSs; (Zhang et al., 2011)). Intuitively, KCIT works by testing for vanishing regression residuals among functions in RKHSs. Another kernel-based CI test called the Permutation Conditional Independence Test (PCIT) reduces CI testing to two-sample kernel-based testing via a carefully chosen permutation found at the solution of a convex optimization problem (Doran et al., 2014).

The aforementioned kernel-based CI tests unfortunately suffer from an important drawback: both tests scale at least quadratically with sample size and therefore take too long to return a p-value in the large sample size setting. In particular, KCIT's bottleneck lies in the eigendecomposition as well as the inversion of large kernel matrices (Zhang et al., 2011), and PCIT takes too long to solve for its required permutation (Doran et al., 2014). As a general rule, it is difficult to develop exact kernel-based methods which scale sub-quadratically with sample size, since the computation of kernel matrices themselves scales at least quadratically.

Many investigators have nonetheless utilized *random Fourier features*, or random features taken from RKHSs, in order to quickly approximate kernel methods. For example, Lopez-Paz and colleagues developed an unconditional independence test using statistics obtained from canonical correlation analysis with random Fourier features (Lopez-Paz et al., 2013). Others have analyzed the use of random Fourier features for predictive modeling (e.g., (Rahimi and Recht, 2007, Sutherland and Schneider, 2015)) or dimensionality reduction (Lopez-

Paz et al., 2014). In practice, investigators have observed that methods which utilize random Fourier features often scale linearly with sample size and achieve comparable accuracy to exact kernel methods.

In this paper, we also use random Fourier features to design two fast tests called the Randomized Conditional Independence Test (RCIT) and the Randomized conditional Correlation Test (RCoT) which approximate the solution of KCIT. Simulations show that RCIT, RCoT and KCIT have comparable accuracy, but both RCIT and RCoT scale linearly with sample size in practice. As a result, RCIT and RCoT return p-values several orders of magnitude faster than KCIT in the large sample size context. Moreover, experiments demonstrate that the causal structures returned by CCD algorithms using either RCIT, RCoT or KCIT have nearly identical accuracy.

2. Characterizations of Conditional Independence

Capital letters X, Y, Z denote sets of random variables with domains $\mathcal{X}, \mathcal{Y}, \mathcal{Z}$, respectively. Consider a measurable, positive definite kernel $k_{\mathcal{X}}$ on \mathcal{X} and denote the corresponding RKHS by $\mathcal{H}_{\mathcal{X}}$. We similarly define $k_{\mathcal{Y}}, \mathcal{H}_{\mathcal{Y}}, k_{\mathcal{Z}}$, and $\mathcal{H}_{\mathcal{Z}}$. We denote the probability distribution of X as \mathbb{P}_X and the joint probability distribution of (X, Z) as \mathbb{P}_{XZ} . Let L^2_X denote the space of square integrable functions of X , and L^2_{XZ} that of (X, Z) . Here, $L^2_X = \{s(X) \mid \mathbb{E}_X(|s|^2) < \infty\}$ and likewise for L^2_{XZ} . Next consider a dataset of n i.i.d. samples drawn according to \mathbb{P}_{XYZ} .

We use the notation $X \perp\!\!\!\perp Y|Z$ when X and Y are conditionally independent given Z . Perhaps the simplest characterization of CI reads as follows: $X \perp\!\!\!\perp Y|Z$ if and only if $\mathbb{P}_{XY|Z} = \mathbb{P}_{X|Z}\mathbb{P}_{Y|Z}$. Equivalently, we have $\mathbb{P}_{X|YZ} = \mathbb{P}_{X|Z}$ and $\mathbb{P}_{Y|XZ} = \mathbb{P}_{Y|Z}$.

2.1. Characterization by RKHSs

A second characterization of CI is given in terms of the cross-covariance operator Σ_{XY} on RKHSs (Fukumizu et al., 2004). For the random vector (X, Y) on $\mathcal{X} \times \mathcal{Y}$, we define the cross-covariance operator from $\mathcal{H}_{\mathcal{Y}}$ to $\mathcal{H}_{\mathcal{X}}$ as follows:

$$\langle f, \Sigma_{XY}g \rangle = \mathbb{E}_{XY}[f(X)g(Y)] - \mathbb{E}_X[f(X)]\mathbb{E}_Y[g(Y)] \quad (1)$$

for all $f \in \mathcal{H}_{\mathcal{X}}$ and $g \in \mathcal{H}_{\mathcal{Y}}$. We may then define the partial cross-covariance operator of (X, Y) given Z by²:

$$\Sigma_{XY \cdot Z} = \Sigma_{XY} - \Sigma_{XZ}\Sigma_{ZZ}^{-1}\Sigma_{ZY}. \quad (2)$$

Notice the similarity of the partial cross-covariance operator to the linear partial cross-covariance matrix (as well as the conditional cross-covariance matrix in the

²Use the right inverse instead of the inverse, if Σ_{ZZ} is not invertible (see Corollary 3 in (Fukumizu et al., 2004)).

Gaussian case)³. Intuitively, one can interpret the above equation as the partial covariance between $\{f(X), \forall f \in \mathcal{H}_X\}$ and $\{g(Y), \forall g \in \mathcal{H}_Y\}$ given $\{h(Z), \forall h \in \mathcal{H}_Z\}$.

Now if we use characteristic kernels⁴ in (2), then the partial cross-covariance operator is related to the CI relation via the following conclusion:

Proposition 1. (Fukumizu et al., 2004, 2008) Let $\ddot{X} = (X, Z)$ and $k_{\ddot{X}} = k_X k_Z$. Also let $\mathcal{H}_{\ddot{X}}$ represent the RKHS corresponding to $k_{\ddot{X}}$. Assume $\mathbb{E}[k_X(X, X)] < \infty$ and $\mathbb{E}[k_Y(Y, Y)] < \infty$ ⁵. Further assume that $k_{\ddot{X}} k_Y$ is a characteristic kernel on $(\mathcal{X} \times \mathcal{Y}) \times \mathcal{Z}$, and that $\mathcal{H}_Z + \mathbb{R}$ (the direct sum of the two RKHSs) is dense in L^2_Z . Then

$$\Sigma_{\ddot{X}Y \cdot Z} = 0 \iff X \perp\!\!\!\perp Y|Z. \quad (3)$$

2.2. Characterization by L^2 spaces

We also consider a different characterization of CI which enforces the uncorrelatedness of functions in suitable spaces; this definition is intuitively more appealing. In particular, consider the following constrained L^2 spaces:

$$\begin{aligned} \mathcal{F}_{XZ} &\triangleq \{f \in L^2_{XZ} \mid \mathbb{E}(f|Z) = 0\}, \\ \mathcal{F}_{YZ} &\triangleq \{g \in L^2_{YZ} \mid \mathbb{E}(g|Z) = 0\}, \\ \mathcal{F}_{Y \cdot Z} &\triangleq \{h \mid h = h'(Y) - \mathbb{E}(h'|Z), h' \in L^2_Y\}. \end{aligned} \quad (4)$$

We then have the following result:

Proposition 2. (Daudin, 1980) The following conditions are equivalent:

1. $X \perp\!\!\!\perp Y|Z$,
2. $\mathbb{E}(fg) = 0, \forall f \in \mathcal{F}_{XZ}$ and $\forall g \in \mathcal{F}_{YZ}$,
3. $\mathbb{E}(fg') = 0, \forall f \in \mathcal{F}_{XZ}$ and $\forall g' \in L^2_{YZ}$,
4. $\mathbb{E}(fh) = 0, \forall f \in \mathcal{F}_{XZ}$ and $\forall h \in \mathcal{F}_{Y \cdot Z}$,
5. $\mathbb{E}(fh') = 0, \forall f \in \mathcal{F}_{XZ}$ and $\forall h' \in L^2_Y$.

The second condition means that any ‘‘residual’’ function of (X, Z) given Z is uncorrelated with that of (Y, Z) given Z . The equivalence also represents a generalization of the case when (X, Y, Z) is jointly Gaussian; here, $X \perp\!\!\!\perp Y|Z$ if

³Recall that the partial cross-covariance of X and Y given Z is defined as $\mathbb{E}[(X - \mathbb{E}(X|Z))(Y - \mathbb{E}(Y|Z))]$; in other words, it is equivalent to the cross-covariance of X and Y given Z . In contrast, the conditional cross-covariance of X and Y given Z is defined as $\mathbb{E}[(X - \mathbb{E}(X|Z))(Y - \mathbb{E}(Y|Z))|Z]$ (notice the extra conditioning).

⁴A kernel k_X is characteristic if $\mathbb{E}_{X \sim \mathbb{P}_X}[f(X)] = \mathbb{E}_{X \sim \mathbb{Q}_X}[f(X)], \forall f \in \mathcal{H}_X$ implies $\mathbb{P}_X = \mathbb{Q}_X$, where \mathbb{P}_X and \mathbb{Q}_X are two probability distributions of X (Fukumizu et al., 2008). Two examples of characteristic kernels include the Gaussian RBF kernel and the Laplacian kernel.

⁵This assumption ensures that $\mathcal{H}_X \subset L^2_X$ and $\mathcal{H}_Y \subset L^2_Y$.

and only if any residual function of X given Z is uncorrelated with that of Y given Z ; i.e., the linear partial correlation coefficient $\rho_{XY.Z}$ is zero.

We also encourage the reader to observe the close relationship between Proposition 1 and claim 4 of Proposition 2. Here, we have almost equivalent statements, but Proposition 1 only considers functions in RKHSs, while claim 4 of Proposition 2 considers functions in L^2 spaces. We find Proposition 1 more useful than claim 4 of Proposition 2 because the RKHS of a characteristic kernel might be much smaller than the corresponding L^2 space.

3. Test Statistic & its Asymptotic Distribution

We consider the following hypotheses:

$$\begin{aligned} H_0 &: X \perp\!\!\!\perp Y|Z, \\ H_1 &: X \not\perp\!\!\!\perp Y|Z. \end{aligned} \tag{5}$$

Now KCIT uses an empirical estimate of the squared Hilbert-Schmidt norm of the partial cross-covariance operator as a statistic to determine whether to reject H_0 :

$$\mathcal{S}_K = n \|\Sigma_{\check{X}Y.Z}\|_{HS}^2. \tag{6}$$

Here, $\|\Sigma_{\check{X}Y.Z}\|_{HS}^2$ denotes an empirical estimate of $\|\Sigma_{\check{X}Y.Z}\|_{HS}^2$, which we can compute using centered kernel matrices (see Theorem 4 and Proposition 5 of (Zhang et al., 2011) for details). We can justify \mathcal{S}_K as a measure of CI due to Proposition 1. We may thus equivalently rewrite the null and alternative in 5 more explicitly as follows:

$$\begin{aligned} H_0 &: \|\Sigma_{\check{X}Y.Z}\|_{HS}^2 = 0, \\ H_1 &: \|\Sigma_{\check{X}Y.Z}\|_{HS}^2 > 0. \end{aligned} \tag{7}$$

In this report, we will also take advantage of the characterization of CI presented in Proposition 1. Recall that the Frobenius norm corresponds to the Hilbert-Schmidt norm in Euclidean space. We therefore consider the squared Frobenius norm of the empirical partial cross-covariance matrix as an approximation of 6 for RCIT:

$$\mathcal{S} = n \|\widehat{\Sigma}_{\check{A}B.C}\|_F^2, \tag{8}$$

where $\widehat{\Sigma}_{\check{A}B.C} = \frac{1}{n-1} \sum_{i=1}^n [(\check{A}_i - \widehat{\mathbb{E}}(\check{A}|C))(B_i - \widehat{\mathbb{E}}(B|C))]$ resembles the empirical cross-covariance matrix. We also have $\check{A} = f'(\check{X}) \triangleq \{f'_1(\check{X}), \dots, f'_m(\check{X})\}$ with $f'_j(\check{X}) \in \mathcal{G}_{\check{X}}, \forall j$. Similarly, $B = h'(Y) \triangleq \{h'_1(Y), \dots, h'_q(Y)\}$ with $h'_k(Y) \in \mathcal{G}_Y, \forall k$, and $C = g(Z) \triangleq \{g_1(Z), \dots, g_d(Z)\}$ with $g_l(Z) \in \mathcal{G}_Z, \forall l$. Here, $\mathcal{G}_{\check{X}}$, \mathcal{G}_Y , and \mathcal{G}_Z denote three spaces of functions, which we will specify shortly. In other words, we select m functions from $\mathcal{G}_{\check{X}}$, q functions from \mathcal{G}_Y , and d functions from \mathcal{G}_Z . We henceforth choose to take the following hypotheses as equivalent to those in 7 and 5:

$$\begin{aligned} H_0 &: \|\Sigma_{\check{A}B.C}\|_F^2 = 0, \\ H_1 &: \|\Sigma_{\check{A}B.C}\|_F^2 > 0. \end{aligned} \tag{9}$$

Now we will compute $\widehat{\Sigma}_{\widehat{A}B.C}$ using $\widehat{\Sigma}_{\widehat{A}B} - \widehat{\Sigma}_{\widehat{A}C}(\widehat{\Sigma}_{CC} + \gamma I)^{-1}\widehat{\Sigma}_{CB}$ similar to 2, where γ denotes a small ridge parameter; recall that this is equivalent to computing the cross-covariance matrix across the residuals of \widehat{A} and B given C using linear ridge regression. We thus may *not* necessarily have $\widehat{\Sigma}_{\widehat{A}B.C} = \widehat{\Sigma}_{\widehat{A}B} - \widehat{\Sigma}_{\widehat{A}C}(\widehat{\Sigma}_{CC} + \gamma I)^{-1}\widehat{\Sigma}_{CB}$. However, we may have $\widehat{\Sigma}_{\widehat{A}B.C} \approx \widehat{\Sigma}_{\widehat{A}B} - \widehat{\Sigma}_{\widehat{A}C}(\widehat{\Sigma}_{CC} + \gamma I)^{-1}\widehat{\Sigma}_{CB}$, if we choose \mathcal{G}_Z in the right way. We therefore must define the space \mathcal{G}_Z in a sensible manner.

In this report, we will set \mathcal{G}_Z to $\{\sqrt{2}\cos(W^T Z + B) | W \sim \mathbb{P}_W, B \sim \text{Uniform}([0, 2\pi])\}$ and likewise for $\mathcal{G}_{\widehat{x}}$ and \mathcal{G}_Y . We select these specific spaces because we can use them to approximate continuous shift-invariant kernels⁶, such as the Gaussian RBF kernel or the Laplacian kernel, via the following result:

Proposition 3. (Rahimi and Recht, 2007) *For a continuous shift-invariant kernel $k(x, y)$ on \mathbb{R}^p , we have:*

$$k(x, y) = \int_{\mathbb{R}^p} e^{iw^T(x-y)} dF_w = \mathbb{E}[\zeta(x)\zeta(y)], \quad (10)$$

where F_W represents the CDF of \mathbb{P}_W and $\zeta(x) = \sqrt{2}\cos(W^T x + B)$ with $W \sim \mathbb{P}_W$ and $B \sim \text{Uniform}([0, 2\pi])$.

The precise form of \mathbb{P}_W depends on the type of shift-invariant kernel one would like to approximate (see Figure 1 of (Rahimi and Recht, 2007) for a list). Since KCIT uses the Gaussian RBF kernel, we choose to approximate the Gaussian RBF kernel by setting \mathbb{P}_W to a Gaussian.

Now let $f_j = f'_j - \mathbb{E}(f'_j|Z)$. Then $\mathbb{E}(f_j|Z) = 0$, so $f_j \in \mathcal{F}_{XZ}$. Moreover, $h'_k - \mathbb{E}(h'_k|Z) \in \mathcal{F}_{Y.Z}$. Note that we can estimate $\mathbb{E}(f'_j|Z)$ with the linear ridge regression solution $\widehat{u}_j^T g(Z)$ under mild conditions because we can guarantee that $\mathbb{P}[|\widehat{\mathbb{E}}_R(f'_j|Z) - \widehat{u}_j^T g(Z)| \geq \varepsilon] \rightarrow 0$ for any fixed $\varepsilon > 0$, where $\widehat{\mathbb{E}}_R(f'_j|Z)$ denotes the estimate of $\mathbb{E}(f'_j|Z)$ by kernel ridge regression; this holds so long as we choose d large enough for $g(Z)$ (see Section 3.1 of (Sutherland and Schneider, 2015); the argument is complex and beyond the scope of this paper). We can also estimate $\mathbb{E}(h'_k|Z)$ with $\widehat{u}_k^T g(Z)$, because we can similarly guarantee that $\mathbb{P}[|\widehat{\mathbb{E}}_R(h'_k|Z) - \widehat{u}_k^T g(Z)| \geq \varepsilon] \rightarrow 0$ for any fixed $\varepsilon > 0$.

We can therefore consider the following spaces for \mathcal{S} which are similar to the L^2 spaces used in claim 4 of Proposition 2:

$$\begin{aligned} \widehat{\mathcal{G}}_{\widehat{x}} &\triangleq \{f \mid f_j = f'_j - \mathbb{E}(f'_j|Z), f'_j \in \mathcal{G}_{\widehat{x}}\}, \\ \widehat{\mathcal{G}}_{Y.Z} &\triangleq \{h \mid h_k = h'_k - \mathbb{E}(h'_k|Z), h'_k \in \mathcal{G}_Y\}. \end{aligned} \quad (11)$$

We then approximate CI with \mathcal{S} in the following sense:

⁶A kernel k is said to be shift-invariant if and only if, for any $a \in \mathbb{R}^p$, we have $k(x-a, y-a) = k(x, y)$, $\forall (x, y) \in \mathbb{R}^p \times \mathbb{R}^p$.

1. We always have $X \perp\!\!\!\perp Y|Z \implies \mathbb{E}(fh) = 0, \forall f \in \widehat{\mathcal{G}}_{\dot{X}}$ and $\forall h \in \widehat{\mathcal{G}}_{Y.Z}$.
2. The reverse direction will hold for an increasing number of distributions as m, q increase.

In practice, we find that the second point holds in all of the cases we tested with only $m, q = 5$.

3.1. Null Distribution

We now consider the asymptotic distribution of \mathcal{S} under the null.

Theorem 1. *Consider n i.i.d. samples from \mathbb{P}_{XYZ} . We then have the following asymptotic distribution under the null in 9:*

$$n\|\widehat{\Sigma}_{\dot{A}B.C}\|_F^2 \xrightarrow{d} \sum_{i=1}^L \lambda_i z_i^2, \quad (12)$$

where $\{z_1, \dots, z_L\}$ denotes i.i.d. standard Gaussian variables (thus $\{z_1^2, \dots, z_L^2\}$ denotes i.i.d. χ_1^2 variables), L the number of elements in $\widehat{\Sigma}_{\dot{A}B.C}$, and λ the eigenvalues of the covariance matrix Π , which we assume to be positive definite; the matrix Π is more specifically the covariance matrix of the vectorization of $(\dot{A} - \mathbb{E}(\dot{A}|C))(B - \mathbb{E}(B|C))^T$. We may denote an arbitrary entry in Π as follows:

$$\begin{aligned} & \Pi_{\dot{A}_i B_j, \dot{A}_k B_l} \\ &= \mathbb{E}[(\dot{A}_i - \mathbb{E}(\dot{A}_i|C))(B_j - \mathbb{E}(B_j|C))(\dot{A}_k - \mathbb{E}(\dot{A}_k|C))(B_l - \mathbb{E}(B_l|C))]. \end{aligned} \quad (13)$$

Proof. We may first write:

$$\begin{aligned} & n\|\widehat{\Sigma}_{\dot{A}B.C}\|_F^2 \\ &= n * \text{tr}(\widehat{\Sigma}_{\dot{A}B.C} \widehat{\Sigma}_{\dot{A}B.C}^T) \\ &= n * v(\widehat{\Sigma}_{\dot{A}B.C})^T v(\widehat{\Sigma}_{\dot{A}B.C}), \\ &= [\sqrt{n}v(\widehat{\Sigma}_{\dot{A}B.C})]^T [\sqrt{n}v(\widehat{\Sigma}_{\dot{A}B.C})], \end{aligned} \quad (14)$$

where $v(\widehat{\Sigma}_{\dot{A}B.C})$ stands for the vectorization of $\widehat{\Sigma}_{\dot{A}B.C}$. By CLT of the sample covariance matrix (see Lemma 1 in the Appendix) combined with the continuous mapping theorem and the null, we know that $\sqrt{n}v(\widehat{\Sigma}_{\dot{A}B.C}) \xrightarrow{d} \mathcal{N}(0, \Pi)$. Here, we write an arbitrary entry $\Pi_{\dot{A}_i B_j, \dot{A}_k B_l}$ under the null as follows:

$$\begin{aligned} & \Pi_{\dot{A}_i B_j, \dot{A}_k B_l} \\ &= \text{Cov}[(\dot{A}_i - \mathbb{E}(\dot{A}_i|C))(B_j - \mathbb{E}(B_j|C)), \\ & \quad (\dot{A}_k - \mathbb{E}(\dot{A}_k|C))(B_l - \mathbb{E}(B_l|C))] \\ &= \mathbb{E}[(\dot{A}_i - \mathbb{E}(\dot{A}_i|C))(B_j - \mathbb{E}(B_j|C))* \\ & \quad (\dot{A}_k - \mathbb{E}(\dot{A}_k|C))(B_l - \mathbb{E}(B_l|C))]. \end{aligned} \quad (15)$$

Now consider the eigendecomposition of Π written as $\Pi = E\Lambda E^T$. Then, we have $E^T[\sqrt{nv}(\widehat{\Sigma}_{\check{A}B.C})] \xrightarrow{d} \mathcal{N}(0, \Lambda)$ by the continuous mapping theorem. Note that:

$$\begin{aligned} & [\sqrt{nv}(\widehat{\Sigma}_{\check{A}B.C})]^T [\sqrt{nv}(\widehat{\Sigma}_{\check{A}B.C})] \\ &= (E^T[\sqrt{nv}(\widehat{\Sigma}_{\check{A}B.C})])^T (E^T[\sqrt{nv}(\widehat{\Sigma}_{\check{A}B.C})]) \\ &\xrightarrow{d} \sum_{i=1}^L \lambda_i z_i^2. \end{aligned} \tag{16}$$

□

We conclude that the null distribution of the test statistic is a positively weighted sum of i.i.d. χ_1^2 random variables. Note that we can obtain estimates of the conditional expectations in Π by using kernel ridge regressions. We will however not need to perform the kernel ridge regressions directly, because we can approximate the outputs of kernel ridge regressions to within an arbitrary degree of accuracy using linear ridge regressions with enough random Fourier features (Sutherland and Schneider, 2015). We can finally obtain an estimate of Π by application of the continuous mapping theorem and the weak law of large numbers. For an arbitrary entry in Π :

$$\begin{aligned} & \frac{1}{n} \sum_{r=1}^n (\check{A}_{i,r} - \widehat{\mathbb{E}}(\check{A}_i|C))(B_{j,r} - \widehat{\mathbb{E}}(B_j|C)) * \\ & \quad (\check{A}_{k,r} - \widehat{\mathbb{E}}(\check{A}_k|C))(B_{l,r} - \widehat{\mathbb{E}}(B_l|C)) \tag{17} \\ & \xrightarrow{P} \mathbb{E}[(\check{A}_i - \mathbb{E}(\check{A}_i|C))(B_j - \mathbb{E}(B_j|C)) * \\ & \quad (\check{A}_k - \mathbb{E}(\check{A}_k|C))(B_l - \mathbb{E}(B_l|C))]. \end{aligned}$$

Unfortunately, a closed form CDF of a positively weighted sum of chi-squared random variables does not exist in general. We can approximate the CDF by Imhof's method which inverts the characteristic function numerically (Imhof, 1961). We should consider Imhof's method as exact, since it provides error bounds and can be used to compute the distribution at a fixed point to within a desired precision (Solomon and Stephens, 1977, Johnson et al., 2002). However, Imhof's method is too computationally intensive for our purposes. We can nonetheless utilize several fast methods which approximate the null by moment matching.

3.2. Approximating the Null Distribution by Moment Matching

We write the cumulants of a positively weighted sum of i.i.d. χ_1^2 random variables as follows:

$$c_r = 2^{r-1}(r-1)! \sum_{i=1}^L \lambda_i^r, \tag{18}$$

where $\lambda = \{\lambda_1, \dots, \lambda_L\}$ denote the weights. We may for example derive the first three cumulants:

$$m_1 = \sum_{i=1}^L \lambda_i, \quad m_2 = 2 \sum_{i=1}^L \lambda_i^2, \quad m_3 = 8 \sum_{i=1}^L \lambda_i^3. \quad (19)$$

We then recover the moments from the cumulants as follows:

$$m_r = c_r + \sum_{i=1}^{r-1} \binom{r-1}{i-1} c_i m_{r-i}, \quad r = 2, 3, \dots \quad (20)$$

Now the Satterthwaite-Welch method (Welch, 1938, Satterthwaite, 1946, Fairfield-Smith, 1936) represents perhaps the simplest and earliest moment matching method. The method matches the first two moments of the sum with a gamma distribution $\Gamma(\hat{g}, \hat{\theta})$. Zhang and colleagues adopted a similar strategy in their paper introducing KCIT (Zhang et al., 2011). Here, we have:

$$\hat{g} = \frac{1}{2} c_1^2 / c_2, \quad \hat{\theta} = c_2 / c_1. \quad (21)$$

We however find the above gamma approximation rather crude. We therefore also consider applying more modern methods to estimating the distribution of a sum of positively weighted chi-squares. Improved methods such as the Hall-Buckley-Eagleson (Hall, 1983, Buckley and Eagleson, 1988) and the Wood F (Wood, 1989) methods match the first three moments of the sum to other distributions in a similar fashion. On the other hand, the Lindsay-Pilla-Basak method (Lindsay et al., 2000) matches the first $2L$ moments to a mixture distribution.

We will focus on the Lindsay-Pilla-Basak method in this paper, since Bodenham & Adams have already determined that the Lindsay-Pilla-Basak method performs the best through extensive experimentation (Bodenham and Adams, 2016, Bodenham, 2015). We therefore choose to use the method as the default method for RCIT. Briefly, the method approximates the CDF under the null $F_{\mathcal{H}_0}$ using a finite mixture of L Gamma CDFs $F_{\Gamma(g, \theta_i)}$:

$$F_{\mathcal{H}_0} = \sum_{i=1}^L \pi_i F_{\Gamma(g, \theta_i)}, \quad (22)$$

where $\pi \geq 0$, $\sum_{i=1}^L \pi_i = 1$, and we seek to determine the $2L + 1$ parameters $g, \theta_1, \dots, \theta_L$, and π_1, \dots, π_L . The Lindsay-Pilla-Basak method computes these parameters by a specific sequence of steps that makes use of results concerning moment matrices (see Appendix II in (Uspensky, 1937)). The sequence is complicated and beyond the scope of this paper, but we refer the reader to (Lindsay et al., 2000) for details.

3.3. Testing for Conditional Un-Correlatedness

Strictly speaking, we must consider the extended variable set \check{X} to test for conditional independence according to Proposition 1. However, we have two observations: (1) we can substitute a test for non-linear conditional uncorrelatedness with tests for conditional independence in almost all cases encountered in practice because most conditionally dependent variables are correlated after some functional transformations, and (2) using the extended variable set \check{X} makes estimating the null distribution more difficult compared to using the unextended variable set X . The first observation coincides with the observations of others who have noticed that the linear partial correlation test performs well (but not perfectly) in ruling out conditional independencies with non-Gaussian data. We can also justify the first observation with the following result using the cross-covariance operator $\Sigma_{XY \cdot Z}$:

Proposition 4. (*Fukumizu et al., 2004, 2008*) Assume $\mathbb{E}[k_X(X, X)] < \infty$ and $\mathbb{E}[k_Y(Y, Y)] < \infty$. Further assume that $k_X k_Y$ is a characteristic kernel on $\mathcal{X} \times \mathcal{Y}$, and that $\mathcal{H}_Z + \mathbb{R}$ (the direct sum of the two RKHSs) is dense in L_Z^2 . Then

$$\Sigma_{XY \cdot Z} = 0 \iff \mathbb{E}_Z[\mathbb{P}_{X|Z}\mathbb{P}_{Y|Z}]. \quad (23)$$

In other words, we have:

$$\begin{aligned} \Sigma_{XY \cdot Z} = 0 &\implies \mathbb{P}_{XY} = \int \mathbb{P}_{X|Z}\mathbb{P}_{Y|Z} d\mathbb{P}_Z, \\ \Sigma_{XY \cdot Z} = 0 &\iff \mathbb{E}_Z[\mathbb{P}_{X|Z}\mathbb{P}_{Y|Z}] = \mathbb{E}_Z[\mathbb{P}_{XY|Z}] \iff X \perp\!\!\!\perp Y|Z. \end{aligned} \quad (24)$$

Notice that $\Sigma_{XY \cdot Z} = 0$ is almost equivalent to CI, in the sense that $\Sigma_{XY \cdot Z} = 0$ just misses those rather contrived distributions where $\mathbb{P}_{XY} = \int \mathbb{P}_{XY|Z} d\mathbb{P}_Z = \int \mathbb{P}_{X|Z}\mathbb{P}_{Y|Z} d\mathbb{P}_Z$ when $X \not\perp\!\!\!\perp Y|Z$. In other words, if $\mathbb{P}_{XY} \neq \int \mathbb{P}_{X|Z}\mathbb{P}_{Y|Z} d\mathbb{P}_Z$ when $X \not\perp\!\!\!\perp Y|Z$, then we have $\Sigma_{XY \cdot Z} = 0 \iff \Sigma_{\check{X}Y \cdot Z} = 0$ (under the corresponding additional assumptions of Propositions 1 and 4).

Let us now consider an example of a situation where $\int \mathbb{P}_{XY|Z} d\mathbb{P}_Z \neq \int \mathbb{P}_{X|Z}\mathbb{P}_{Y|Z} d\mathbb{P}_Z$ when $X \not\perp\!\!\!\perp Y|Z$. Take three binary variables $X, Y, Z \in \{0, 1\}$. Let $\mathbb{P}_{Z=0} = 0.2$ and $\mathbb{P}_{Z=1} = 0.8$. Also consider the four probability tables in Table 1. Here, we have chosen the probabilities in the tables carefully by satisfying the following equation:

$$\begin{aligned} \int \mathbb{P}_{XY|Z} d\mathbb{P}_Z &= \int \mathbb{P}_{X|Z}\mathbb{P}_{Y|Z} d\mathbb{P}_Z \\ \iff \mathbb{P}_{Z=0}(\mathbb{P}_{X|Z=0}\mathbb{P}_{Y|Z=0}) + \mathbb{P}_{Z=1}(\mathbb{P}_{X|Z=1}\mathbb{P}_{Y|Z=1}) & \\ = \mathbb{P}_{Z=0}\mathbb{P}_{XY|Z=0} + \mathbb{P}_{Z=1}\mathbb{P}_{XY|Z=1}. & \end{aligned} \quad (25)$$

Of course, the equality holds when we have conditional independence $\mathbb{P}_{XY|Z} = \mathbb{P}_{X|Z}\mathbb{P}_{Y|Z}$. We are however interested in the case when conditional dependence holds. We therefore instantiated the values of Tables 1a and 1b as well as the second column in Table 1c ($\mathbb{P}_{XY|Z=0}$) such that $\mathbb{P}_{XY|Z=0} \neq \mathbb{P}_{X|Z=0}\mathbb{P}_{Y|Z=0}$. We

	$\mathbb{P}_{X Z=0}$	$\mathbb{P}_{X Z=1}$
$X = 0$	0.5	0.3
$X = 1$	0.5	0.7

(a)

	$\mathbb{P}_{Y Z=0}$	$\mathbb{P}_{Y Z=1}$
$Y = 0$	0.3	0.4
$Y = 1$	0.7	0.6

(b)

	$\mathbb{P}_{XY Z=0}$	$\mathbb{P}_{XY Z=1}$
$X = 0, Y = 0$	0.2	0.1075
$X = 0, Y = 1$	0.3	0.1925
$X = 1, Y = 0$	0.1	0.2925
$X = 1, Y = 1$	0.4	0.4075

(c)

	\mathbb{P}_{XY}
$X = 0, Y = 0$	0.126
$X = 0, Y = 1$	0.214
$X = 1, Y = 0$	0.254
$X = 1, Y = 1$	0.406

(d)

Table 1: Example of a situation where $\int \mathbb{P}_{XY|Z} d\mathbb{P}_Z = \int \mathbb{P}_{X|Z} \mathbb{P}_{Y|Z} d\mathbb{P}_Z$ when $X \not\perp\!\!\!\perp Y|Z$ using binary variables.

then solved for $\mathbb{P}_{XY|Z=1}$ using Equation 25 in order to complete Table 1c. This ultimately yielded Table 1d.

Notice that we obtain a unique value for $\mathbb{P}_{XY|Z=1}$ by solving Equation 25. Hence, $\mathbb{P}_{XY|Z=1}$ has Lebesgue measure zero on the interval $[0, 1]$, once we have defined all of the other variables in the equation. Thus, $\Sigma_{XY \cdot Z} = 0$ is not always equivalent to $X \perp\!\!\!\perp Y|Z$, but satisfying the condition $\int \mathbb{P}_{XY|Z} d\mathbb{P}_Z = \int \mathbb{P}_{X|Z} \mathbb{P}_{Y|Z} d\mathbb{P}_Z$ when $X \not\perp\!\!\!\perp Y|Z$ requires a very particular setup which is probably rarely encountered in practice.

The aforementioned argument motivates us to also consider the following statistic using a finite dimensional partial cross-covariance matrix:

$$\mathcal{S}' = n \|\widehat{\Sigma}_{AB \cdot C}\|_F^2, \quad (26)$$

where we have replaced \check{A} with A . The above statistic is a generalization of linear partial correlation, because we consider uncorrelatedness of the residuals of non-linear functional transformations after performing non-linear regression. The asymptotic distribution for \mathcal{S} in Theorem 1 also holds for \mathcal{S}' , when we replace \check{A} with A . Here, we use the hypotheses:

$$\begin{aligned} H_0 &: \|\Sigma_{AB \cdot C}\|_F^2 = 0, \\ H_1 &: \|\Sigma_{AB \cdot C}\|_F^2 > 0. \end{aligned} \quad (27)$$

In practice, the test which uses \mathcal{S}' , which we now call the Randomized conditional Correlation Test (RCoT), usually rivals or outperforms RCIT and KCIT, because (1) nearly all conditionally dependent variables encountered in practice are also conditionally correlated after at least one functional transformation, and (2) we can easily calibrate the null distribution of the test using \mathcal{S}' even when Z has large cardinality. We will therefore find this test useful for replacing RCIT when we have large conditioning set sizes (≥ 4).

4. Experiments

We carried out experiments to compare the empirical performance of the following tests:

- RCIT: uses \mathcal{S} with the Lindsay-Pilla-Basak approximation,
- RCoT: uses \mathcal{S}' with the Lindsay-Pilla-Basak approximation,
- KCIT: uses \mathcal{S}_K with a simulated null by bootstrap.

Note that KCIT with the gamma approximation performs *slightly* faster than KCIT with bootstrap (e.g., less than 200ms faster on average at 2000 samples in our experiments), but the bootstrap results in a significantly better calibrated null distribution. We focus on large sample size (≥ 500) scenarios because we can just apply KCIT with bootstrap otherwise. We ran all experiments using the R programming language (Microsoft R Open) on a laptop with 2.60 GHz of CPU and 16GB of RAM.

4.1. Hyperparameters

We used the same hyperparameters for RCIT and RCoT. Namely, we used the median Euclidean distance heuristic across the first 500 samples of \check{X} , X , Y and Z for choosing the $\sigma_{\check{X}}$, σ_X , σ_Y , and σ_Z hyperparameters for the Gaussian RBF kernels $k_\sigma(x, y) = \exp(-\|x - y\|^2/\sigma)$, respectively⁷ (Gretton et al., 2008, Lopez-Paz et al., 2014). We also fixed the number of Fourier features for \check{X} , X and Y to 5 and the number of Fourier features for Z to 25. We standardized all original and Fourier variables to mean zero unit variance in order to help ensure numerically stable computations. Finally, we set γ to 1E-10 in order to keep bias minimal.

With KCIT, we set σ to the squared median Euclidean distance between (X, Y) using the first 500 samples times double the conditioning set size; the hyperparameters as described in the original paper, the hyperparameters in the author-provided MATLAB implementation and the hyperparameters of RCIT/RCoT all gave worse performance.

4.2. Type I Error

A good statistical test should control the Type I error rate at any specified α . We therefore analyzed the Type I error rates of the three CI tests as a function of sample size and conditioning set size. We evaluated the algorithms using the Kolmogorov-Smirnov (KS) test statistic. Recall that the KS test uses the following statistic:

$$\mathcal{K} = \sup_{x \in \mathbb{R}} |\widehat{F}(x) - F(x)| = \|\widehat{F}_X - F_X\|_\infty, \quad (28)$$

⁷We also tried setting σ_Z to the median distance divided by 1.5, 2 or 3. However, these values gave progressively worse performance on average.

where \widehat{F}_X denotes the empirical CDF, and F_X some comparison CDF. If the sample comes from \mathbb{P}_X , then \mathcal{K} converges to 0 almost surely as $n \rightarrow \infty$ by the Glivenko-Cantelli theorem.

Now a good CI test controls the Type I error rate at any α value, when we have a uniform sampling distribution of the p-values over $[0, 1]$. Therefore, a good CI test should have a small KS statistic value, when we set F_X to the uniform distribution over $[0, 1]$.

To compute the KS statistic values, we generated data from 1000 post non-linear models (Zhang et al., 2011, Doran et al., 2014). We can describe each post non-linear model as follows: $X = g_1(Z + \varepsilon_1), Y = g_2(Z + \varepsilon_2)$, where $Z, \varepsilon_1, \varepsilon_2$ have jointly independent standard Gaussian distributions, and g_1, g_2 denote smooth functions. We always chose g_1, g_2 uniformly from the following set of functions: $\{(\cdot), (\cdot)^2, (\cdot)^3, \tanh(\cdot), \exp(-\|\cdot\|_2)\}$. Thus, we have $X \perp\!\!\!\perp Y|Z$ in any case. Notice also that this situation is more general than the additive noise models proposed in (Ramsey, 2014), where we have $X = g_1(Z) + \varepsilon_1, Y = g_2(Z) + \varepsilon_2$.

4.2.1. Sample Size

We first assess the Type I error rate as a function of sample size. We used sample sizes of 500, 1000, 2000, 5000, ten thousand, one hundred thousand and one million. A good CI test should control the Type I error rate across all α values at any sample size. Figure 1a summarizes the KS statistic values for the three different CI tests. Observe that all tests have similar KS statistic values across different sample sizes. We conclude that all three tests perform comparably in controlling the Type I error rate with a single conditioning variable at different sample sizes.

The run time results however tell a markedly different story. Both RCIT and RCoT output a p-value much more quickly than KCIT at different sample sizes (Figure 1b). Moreover, KCIT ran out of memory at 5000 samples while RCIT and RCoT handled one million samples in a little over 6 seconds. RCIT and RCoT also completed more than two orders of magnitude faster than KCIT on average at a sample size of 2000 (Figure 1c). We conclude that RCIT and RCoT are more scalable than KCIT. Moreover, the experimental results agree with standard matrix complexity theory; RCIT and RCoT scale linearly with sample size, while KCIT scales cubically with sample size.

4.2.2. Conditioning Set Size

CCD algorithms request p-values from CI tests using large conditioning set sizes. In fact, algorithms which do not assume causal sufficiency, such as FCI, often demand very large conditioning set sizes (> 5). We should however also realize that CCD algorithms search for *minimal* conditioning sets in order to establish ancestral relations. This means that we must focus on testing for cases where $X \not\perp\!\!\!\perp Y|Z$, but we have either $X \perp\!\!\!\perp Y|Z \cup A$ or $X \not\perp\!\!\!\perp Y|Z \cup A$, where $|A| = 1$.

We therefore evaluated the Type I error rates of the CI tests as a function of conditioning set size by fixing the sample size at 1000 and then adding 1 to 10 standard Gaussian variables into the conditioning set so that $X =$

$g_1(\frac{1}{k} \sum_{j=1}^k Z_j + \epsilon_1), Y = g_2(\frac{1}{k} \sum_{j=1}^k Z_j + \epsilon_2), k = \{1, \dots, 10\}$ in 1000 models. Note that this situation corresponds to 1 to 10 common causes.

Figure 1d summarizes the KS statistic values in the aforementioned scenario. We see that the KS statistic values for RCoT remain the smallest for nearly all conditioning set sizes, followed by RCIT and then KCIT. This implies that RCoT best approximates the null distribution out of the three CI tests. We also provide the histograms of the p-values across the 1000 post non-linear models at a conditioning set size of 10 for KCIT, RCIT, and RCoT in Figures 1e-1g. Notice that the histograms become progressively more similar to a uniform distribution. We conclude that RCoT controls its Type I error rate the best even with large conditioning set sizes while KCIT controls its rate the worst.

Now the run times of all three tests only increased very slightly with the conditioning set size (Figure 1h). However, both RCIT and RCoT still completed 40.91 times faster than KCIT on average (95% confidence interval: ± 0.44). These results agree with standard matrix complexity theory, as we expect all tests to scale linearly with dimensionality.

4.3. Power

We next evaluated test power (i.e., $1 - (\text{Type II error rate})$) by computing the area under the power curve (AUPC). The AUPC corresponds to the area under the empirical CDF of the p-values returned by a CI test when the null does not hold. A CI test has higher power when its AUPC is closer to one. For example, observe that if a CI test always returns a p-value of 0 in the perfect case, then its AUPC corresponds to 1.

We examined the AUPC by adding the same small error $\epsilon_b \sim \mathcal{N}(0, 1/16)$ to both X and Y in 1000 post non-linear models as follows: $X = g_1(\epsilon_b + \epsilon_1), Y = g_2(\epsilon_b + \epsilon_2), Z \sim \mathcal{N}(0, 1)$. Here, we do not allow the CI tests to condition on ϵ_b , so we always have $X \not\perp Y|Z$; this situation therefore corresponds to a hidden common cause.

4.3.1. Sample Size

We first examine power as a function of sample size. We again tested sample sizes of 500, 1000, 2000, 5000, ten thousand, one hundred thousand, and one million. We have summarized the results in Figure 2a. Both RCIT and RCoT have comparable AUPC values to KCIT with sample sizes of 500, 1000 and 2000. At larger sample sizes, KCIT again did not scale due to insufficient memory, but the AUPC of both RCIT and RCoT continued to increase at similar values. We conclude that all three tests have similar power.

The run time results mimic those of Section 4.2.1; RCIT and RCoT completed orders of magnitude faster than KCIT (Figures 2b and 2c).

4.3.2. Dimensionality

We next examined power as a function of conditioning set size. To do this, we fixed the sample size at 1000 and set $Z = (Z_1, \dots, Z_k)$ with $Z \sim \mathcal{N}(0, I_k), k = \{1, \dots, 10\}$ in the 1000 post non-linear models. We therefore examined how

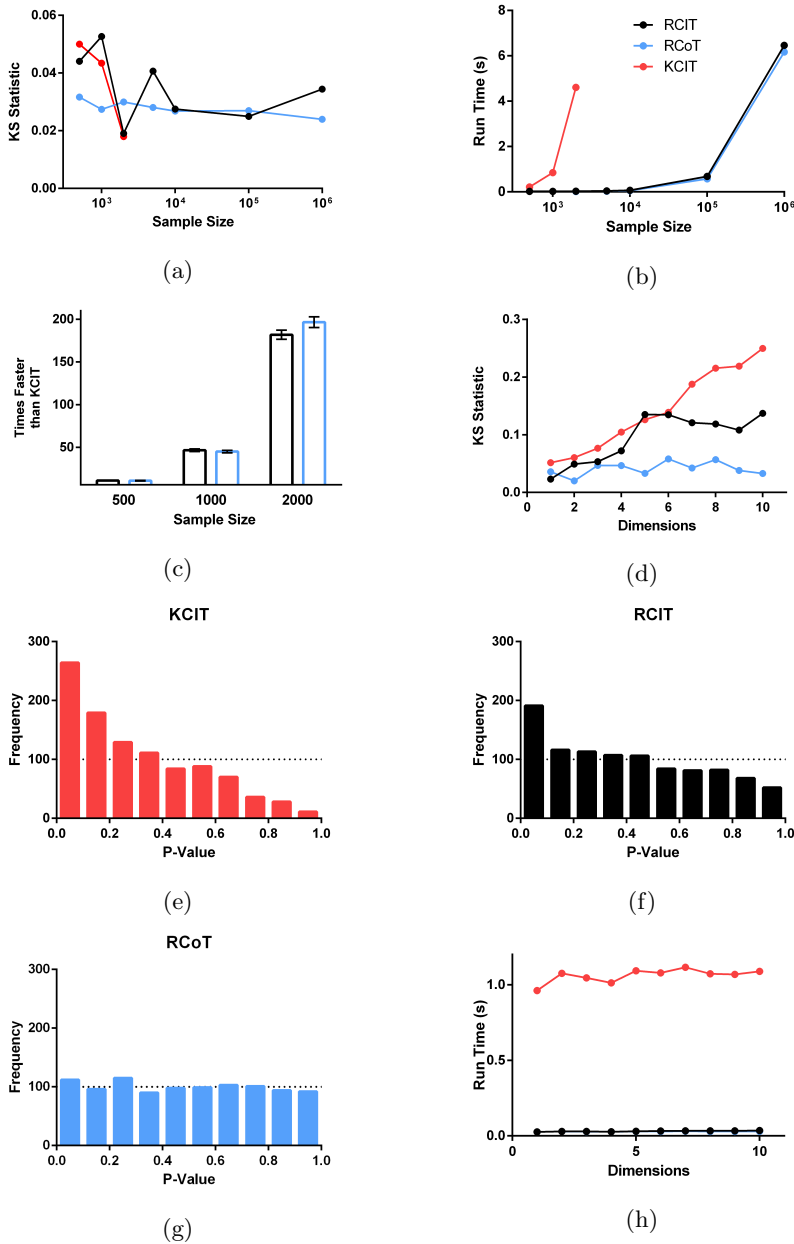


Figure 1: Experimental results of RCIT, RCoT and KCIT as a function of sample size and conditioning set size when conditional independence holds. (a) All tests have comparable KS statistic values as a function of sample size with a conditioning set size of one. (b) However, both RCIT and RCoT complete much faster than KCIT. (c) The relative difference in speed between RCIT vs. KCIT and RCoT vs. KCIT quickly grows with increasing sample size. (d) RCoT maintains the lowest KS statistic value with increases in the dimensionality of the conditioning set. (e-g) Histograms with a conditioning set size of 10 (the hardest case tested) show that KCIT, RCIT and RCoT obtain progressively more uniform null distributions. (h) Run times of all three tests scale linearly with dimensionality of the conditioning set. Both RCIT and RCoT have nearly identical run times in this case, so the black and blue lines overlap.

well the CI tests reject the null under an increasing conditioning set size with uninformative variables. A good CI test should either (1) maintain its power or, more realistically, (2) suffer a graceful decline in power with an increasing conditioning set size because none of the variables in the conditioning set are informative for rendering conditional independence by design.

We have summarized the results in Figure 2d. Notice that all tests have comparable AUPC values with small conditioning set sizes (between 1 and 3), but the AUPC value of KCIT gradually increases with increasing conditioning set sizes; the AUPC value should not increase under the current setup with a well-calibrated null because the extra variables are uninformative. To determine the cause of the unexpected increase in power, we permuted the values of X in each run in order to assess the calibration of the null distribution. Figure 2f summarizes the results, where we can see that only KCIT’s KS statistic grows with an increasing conditioning set size. We can therefore claim that the increasing AUPC value of KCIT holds because of a badly calibrated null distribution with larger conditioning set sizes. We conclude that both RCIT and RCoT maintain steady power under an increasing conditioning set size while KCIT does not.

The run times in Figures 2e and 2g again mimic those in Section 4.2.2 with RCIT and RCoT completing in a much shorter time frame than KCIT.

4.4. Causal Structure Discovery

We next examine the accuracy of graphical structures as recovered by PC (Spirtes et al., 2000), FCI (Zhang, 2008) and RFCI (Colombo et al., 2012) when run using RCIT, RCoT or KCIT.

We used the following procedure in (Colombo et al., 2012) to generate 250 different Gaussian DAGs with an expected neighborhood size $\mathbb{E}(N) = 2$ and $v = 20$ vertices. First, we generated a random adjacency matrix \mathcal{A} with independent realizations of $\text{Bernoulli}(\mathbb{E}(N)/(v-1))$ random variables in the lower triangle of the matrix and zeroes in the remaining entries. Next, we replaced the ones in \mathcal{A} by independent realizations of a $\text{Uniform}([-1, -0.1] \cup [0.1, 1])$ random variable. We interpret a nonzero entry \mathcal{A}_{ij} as an edge from X_i to X_j with coefficient \mathcal{A}_{ij} in the following linear model:

$$\begin{aligned} X_1 &= \varepsilon_1, \\ X_i &= \sum_{r=1}^{v-1} \mathcal{A}_{ir} X_r + \varepsilon_i. \end{aligned} \tag{29}$$

for $i = 2, \dots, v$ where $\varepsilon_1, \dots, \varepsilon_v$ are mutually independent standard Gaussian random variables. The variables $\{X_1, \dots, X_v\} = \mathbf{X}$ then have a multivariate Gaussian distribution with mean 0 and covariance matrix $\Sigma = (I_v - \mathcal{A})^{-1}(I_v - \mathcal{A})^{-T}$, where I_v is the $v \times v$ identity matrix. To introduce non-linearities, we passed each variable in \mathbf{X} through a non-linear function g again chosen uniformly from the set $\{(\cdot), (\cdot)^2, (\cdot)^3, \tanh(\cdot), \exp(-\|\cdot\|_2)\}$.

For FCI and RFCI, we introduced latent and selection variables using the following procedure. For each DAG, we first randomly selected a set of 0-3




Figure 2: Experimental results with RCIT, RCoT and KCIT as a function of sample size and conditioning set size when conditional dependence holds. (a) All tests have comparable AUPC values as a function of sample size with a conditioning set size of one. (b-c) Both RCIT and RCoT again complete much faster than KCIT. (d) KCIT's AUPC value unexpectedly increases with the dimensionality of the conditioning set. Associated run times for (d) in (e). (f) The cause of KCIT's AUPC increase lies in a badly calibrated null distribution; here we see that only KCIT's KS statistic value increases under the null. Associated run times for (f) in (g).

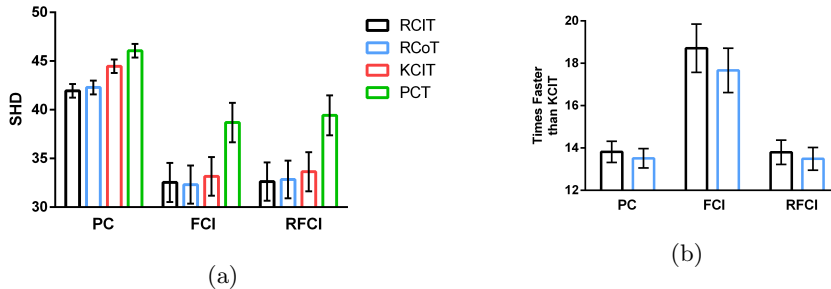


Figure 3: Results of CCD algorithms as evaluated by mean (a) SHD and (b) run times. The CCD algorithms run with KCIT perform comparably (or even slightly worse) to those run with RCIT and RCoT in (a). Run times in (b) show that the CCD algorithms run with RCIT and RCoT complete at least 13 times faster on average than those with KCIT. Error bars denote 95% confidence intervals of the mean.

latent common causes L . From the set $X \setminus L$, we then selected a set of 0-3 colliders as selection variables S . For each selection variable in S , we subsequently eliminated the bottom q percentile of samples, where we drew q according to independent realizations of a $\text{Uniform}([0.1, 0.5])$ random variable. We finally eliminated all of the latent variables from the dataset.

We ultimately created 250 different 500 sample datasets for PC, FCI and RFCI. We then ran the sample versions of PC, FCI and RFCI using RCIT, RCoT, KCIT and the linear partial correlation test (PCT) at $\alpha = 0.05$. We also obtained the oracle graphs by running the oracle versions of PC, FCI and RFCI using the ground truth.

We have summarized the results as structural Hamming distances (SHDs) from the oracle graphs in Figure 3a. PC run with RCIT and PC run with RCoT both outperformed PC run with KCIT by a large margin according to paired t-tests (PC RCIT vs. KCIT, $t = -14.76, p < 2.2\text{E-}16$; PC RCoT vs. KCIT, $t = -12.87, p < 2.2\text{E-}16$). We found similar results with FCI and RFCI, although by only a small margin; 3 of the 4 comparisons fell below the Bonferonni corrected threshold of $0.05/6$ and the other comparison fell below the uncorrected threshold of 0.05 (FCI RCIT vs. KCIT, $t = -2.00, p = 0.047$; FCI RCoT vs. KCIT $t = -2.96, p = 0.0034$; RFCI RCIT vs. KCIT, $t = -3.56, p = 4.5\text{E-}4$; RFCI RCoT vs. KCIT, $t = -2.80, p = 0.0055$). All algorithms with any of the kernel-based tests outperformed the same algorithms with PCT by a large margin ($p < 7\text{E-}14$ in all cases). Finally, the run time results in Figure 3b show that the CCD algorithms run with RCIT and RCoT complete at least 13 times faster on average than those run with KCIT. We conclude that both RCIT and RCoT help CCD algorithms at least match the performance of the same algorithms run with KCIT, but RCIT and RCoT do so within a much shorter time frame than KCIT.

4.5. Real Data

We finally ran PC, FCI and RFCI using RCIT, RCoT, KCIT and PCT at $\alpha = 0.05$ on a publicly available longitudinal dataset from the Cognition and Aging USA (CogUSA) study (McArdle et al., 2015), where scientists measured the cognition of men and women above 50 years of age. The dataset contains 815 samples, 18 variables and two waves (thus $18/2 = 9$ variables in each wave) separated by two years after some data cleaning⁸. Note that we do not have access to a gold standard solution set in this case. However, we can utilize the time information in the dataset to detect false positive ancestral relations directed backwards in time.

We ran the CCD algorithms on 30 bootstrapped datasets. Results are summarized in Figure 4. Comparisons with PC did not reach the Bonferroni level among the kernel-based tests, although PC run with either RCIT or RCoT yielded fewer false positive ancestral relations on average than PC run with KCIT near an uncorrected level of 0.05 (PC RCIT vs. KCIT, $t = -2.76, p = 9.85E-3$; PC RCoT vs. KCIT, $t = -1.99, p = 0.056$). However, FCI and RFCI run with either RCIT or RCoT performed better than those run with KCIT at a Bonferroni corrected level of $0.05/6$ (FCI RCIT vs. KCIT, $t = -16.03, p = 6.00E-16$; FCI RCoT vs. KCIT, $t = -9.62, p = 1.59E-10$; RFCI RCIT vs. KCIT, $t = -4.37, p = 1.44E-4$; RFCI RCoT vs. KCIT, $t = -4.57, p = 8.33E-5$). The CCD algorithms run with PCT also gave inconsistent results; PC run with PCT performed the best on average, but FCI and RFCI run with PCT also performed the worst. Here, we should trust the outputs of FCI and RFCI more strongly than those of PC, since both FCI and RFCI allow latent common causes and selection bias which often exist in real data. Next, CCD algorithms run with RCIT performed comparably to those run with RCoT (PC RCIT vs. RCoT, $t = -1.05, p = 0.301$; FCI RCIT vs. RCoT, $t = -2.13, p = 0.042$; RFCI RCIT vs. RCoT, $t = -0.94, p = 0.355$). We finally report that the CCD algorithms run with RCIT and RCoT complete at least 40 times faster on average than those run with KCIT (Figure 4b). We conclude that CCD algorithms run with either RCIT or RCoT perform at least as well as those run with KCIT on this real dataset but with large reductions run time.

5. Conclusion

We developed two statistical tests called RCIT and RCoT for fast non-parametric CI testing. Both RCIT and RCoT approximate KCIT by sampling Fourier features. Moreover, the proposed tests return p-values orders of magnitude faster than KCIT in the large sample size setting. RCoT in particular also has a better calibrated null distribution than KCIT especially with larger conditioning set sizes. In causal graph discovery, RCIT and RCoT help CCD

⁸We specifically removed redundant variables with deterministic relations, variables with more than 1000 missing values, and then samples with missing values in any of the remaining variables.




Figure 4: Results of CCD algorithms as evaluated on real longitudinal data. Part (a) displays mean counts of the number of ancestral relations directed backwards time. We do not display 95% confidence intervals when we computed a standard error of zero. Part (b) summarizes the mean run times.

algorithms recover graphical structures at least as accurately as KCIT but, most importantly, also allow the algorithms to complete in a much shorter time frame. We believe that the speedups provided by RCIT and RCoT will make non-parametric causal discovery more accessible to scientists who wish to apply CCD algorithms to their datasets.

Acknowledgments

Research reported in this publication was supported by grant U54HG008540 awarded by the National Human Genome Research Institute through funds provided by the trans-NIH Big Data to Knowledge initiative. The research was also supported by the National Library of Medicine of the National Institutes of Health under award numbers T15LM007059 and R01LM012095. The content is solely the responsibility of the authors and does not necessarily represent the official views of the National Institutes of Health.

References

References

- D. Bodenham. *momentchi2*. 2015. URL <http://cran.r-project.org/web/packages/momentchi2/>.
- D. Bodenham and N. Adams. A comparison of efficient approximations for a weighted sum of chi-squared random variables. *Statistics and Computing*, 26: 917–928, 2016. doi: 10.1007/s11222-015-9583-4. URL <http://dx.doi.org/10.1007/s11222-015-9583-4>.
- M. J. Buckley and G. K. Eagleson. An approximation to the distribution of quadratic forms in normal random variables. *Australian and New Zealand Journal of Statistics*, 30(1):150–159, 1988.

- D. Colombo, M. Maathius, M. Kalisch, and T. Richardson. Learning high-dimensional directed acyclic graphs with latent and selection variables. *Annals of Statistics*, 40(1):294–321, Apr. 2012. doi: 10.1214/11-AOS940. URL <http://projecteuclid.org/euclid.aos/1333567191>.
- J. J. Daudin. Partial association measures and an application to qualitative regression. 67(3):581–590, Dec. 1980. ISSN 0006-3444 (print), 1464-3510 (electronic). doi: <http://dx.doi.org/10.1093/biomet/67.3.581>; <http://dx.doi.org/10.2307/2335127>. URL <http://www.jstor.org/stable/2335127>.
- G. Doran, K. Muandet, K. Zhang, and B. Schölkopf. A permutation-based kernel conditional independence test. In *Proceedings of the 30th Conference on Uncertainty in Artificial Intelligence (UAI2014)*, pages 132–141, Oregon, 2014. AUAI Press Corvallis.
- H. Fairfield-Smith. The problem of comparing the results of two experiments with unequal errors. *Journal of the Council for Scientific and Industrial Research*, 9:211–212, 1936.
- R. A. Fisher. Frequency distribution of the values of the correlation coefficient in samples from an indefinitely large population. *Biometrika*, 10(4):507–521, 1915. ISSN 00063444. doi: 10.2307/2331838. URL <http://dx.doi.org/10.2307/2331838>.
- R. A. Fisher. On the probable error of a coefficient of correlation deduced from a small sample. *Metron*, 1:3–32, 1921.
- K. Fukumizu, F. R. Bach, and M. I. Jordan. Dimensionality reduction for supervised learning with reproducing kernel hilbert spaces. *J. Mach. Learn. Res.*, 5:73–99, Dec. 2004. ISSN 1532-4435. URL <http://dl.acm.org/citation.cfm?id=1005332.1005335>.
- K. Fukumizu, A. Gretton, X. Sun, and B. Schölkopf. Kernel measures of conditional dependence. In *Advances in Neural Information Processing Systems*, pages 489–496, Red Hook, NY, USA, Sept. 2008. Max-Planck-Gesellschaft, Curran. URL <http://papers.nips.cc/paper/3340-kernel-measures-of-conditional-dependence-supplemental.zip>.
- A. Gretton, K. Fukumizu, C. Teo, L. Song, B. Schölkopf, and A. Smola. A kernel statistical test of independence. In *Advances in neural information processing systems 20*, pages 585–592, Red Hook, NY, USA, Sept. 2008. Max-Planck-Gesellschaft, Curran.
- P. Hall. Chi squared approximations to the distribution of a sum of independent random variables. *Ann. Probab.*, 11(4):1028–1036, 11 1983. doi: 10.1214/aop/1176993451. URL <http://dx.doi.org/10.1214/aop/1176993451>.
- T.-M. Huang. Testing conditional independence using maximal nonlinear conditional correlation. *The Annals of Statistics*, 38(4):2047–2091, 2010.

- J. P. Imhof. Computing the distribution of quadratic forms in normal variables. *Biometrika*, 48(3/4):419–426, 1961. doi: 10.2307/2332763. URL <http://dx.doi.org/10.2307/2332763>.
- N. L. Johnson, S. Kotz, and N. Balakrishnan. *Continuous Multivariate Distributions*. Wiley, 3rd edition, 2002.
- B. Lindsay, R. Pilla, and P. Basak. Moment-based approximations of distributions using mixtures: Theory and applications. *Annals of the Institute of Statistical Mathematics*, 52(2):215–230, 2000. URL <http://EconPapers.repec.org/RePEc:spr:aistmt:v:52:y:2000:i:2:p:215-230>.
- D. Lopez-Paz, P. Hennig, and B. Schölkopf. The randomized dependence coefficient. In *Advances in Neural Information Processing Systems 26*, pages 1–9, 2013.
- D. Lopez-Paz, S. Sra, A. Smola, Z. Ghahramani, and B. Schölkopf. Randomized nonlinear component analysis. In *Proceedings of the 31st International Conference on Machine Learning, W&CP 32 (1)*, pages 1359–1367. JMLR, 2014.
- D. Margaritis. Distribution-free learning of bayesian network structure in continuous domains. In *Proceedings, The Twentieth National Conference on Artificial Intelligence and the Seventeenth Innovative Applications of Artificial Intelligence Conference, July 9-13, 2005, Pittsburgh, Pennsylvania, USA*, pages 825–830, 2005. URL <http://www.aaai.org/Library/AAAI/2005/aaai05-130.php>.
- J. McArdle, W. Rodgers, and R. Willis. Cognition and aging in the usa (cogusa), 2007-2009, 2015.
- K. Pearson. On the criterion that a given system of deviations from the probable in the case of a correlated system of variables is such that it can be reasonably supposed to have arisen from random sampling. *Philosophical Magazine Series 5*, 50:157–175, 1900.
- A. Rahimi and B. Recht. Random features for large-scale kernel machines. In *In Neural Information Processing Systems*, 2007.
- J. D. Ramsey. A scalable conditional independence test for nonlinear, non-gaussian data. *CoRR*, abs/1401.5031, 2014. URL <http://arxiv.org/abs/1401.5031>.
- F. Satterthwaite. : An approximate distribution of estimates of variance components. *Biom. Bull.*, 2(6):110–114, 1946.
- H. Solomon and M. A. Stephens. Distribution of a sum of weighted chi-square variables. *Journal of the American Statistical Association*, 72:881–885, 1977.

- P. Spirtes, C. Glymour, and R. Scheines. *Causation, Prediction, and Search*. MIT press, 2nd edition, 2000.
- L. Su and H. White. A consistent characteristic function-based test for conditional independence. *Journal of Econometrics*, 141(2): 807 – 834, 2007. ISSN 0304-4076. doi: DOI:10.1016/j.jeconom.2006.11.006. URL <http://www.sciencedirect.com/science/article/B6VCO-4MT59DD-4/2/267e7fc8dd979b6148fc4123998e94ee>.
- L. Su and H. White. A nonparametric hellinger metric test for conditional independence. *Econometric Theory*, 24(4):829–864, 2008. ISSN 02664666, 14694360. URL <http://www.jstor.org/stable/20142523>.
- D. J. Sutherland and J. G. Schneider. On the error of random fourier features. In M. Meila and T. Heskes, editors, *UAI*, pages 862–871. AUAI Press, 2015. ISBN 978-0-9966431-0-8. URL <http://dblp.uni-trier.de/db/conf/uai/uai2015.html#SutherlandS15>.
- I. Tsamardinos and G. Borboudakis. *Permutation Testing Improves Bayesian Network Learning*, pages 322–337. Springer Berlin Heidelberg, Berlin, Heidelberg, 2010. ISBN 978-3-642-15939-8. doi: 10.1007/978-3-642-15939-8_21. URL http://dx.doi.org/10.1007/978-3-642-15939-8_21.
- J. V. J. V. Uspensky. *Introduction to mathematical probability*. New York ; London : McGraw-Hill, 1st ed edition, 1937. "Problems for solution" with answers at end of each chapter.
- B. L. Welch. The significance of the difference between two means when the population variances are unequal. *Biometrika*, 29(3-4):350–362, Feb. 1938. doi: 10.1093/biomet/29.3-4.350. URL <http://dx.doi.org/10.1093/biomet/29.3-4.350>.
- A. T. A. Wood. An f approximation to the distribution of a linear combination of chi-squared variables. *Communications in Statistics: Simulation and Computation*, 18:1439–1456, 1989.
- J. Zhang. On the completeness of orientation rules for causal discovery in the presence of latent confounders and selection bias. *Artif. Intell.*, 172(16-17): 1873–1896, Nov. 2008. ISSN 0004-3702. doi: 10.1016/j.artint.2008.08.001. URL <http://dx.doi.org/10.1016/j.artint.2008.08.001>.
- K. Zhang, J. Peters, D. Janzing, and B. Schölkopf. Kernel-based conditional independence test and application in causal discovery. In *Uncertainty in Artificial Intelligence*, pages 804–813. AUAI Press, 2011. ISBN 978-0-9749039-7-2. URL <http://dblp.uni-trier.de/db/conf/uai/uai2011.html#ZhangPJS11>.

6. Appendix

We will prove the central limit theorem (CLT) for the sample covariance matrix. We first have the following sample covariance matrices with known and unknown expectation vector, respectively:

$$\begin{aligned}\ddot{\Sigma} &= \frac{1}{n} \sum_{i=1}^n [X_i - \mathbb{E}(X)] [X_i - \mathbb{E}(X)]^T, \\ \widehat{\Sigma} &= \frac{1}{n-1} \sum_{i=1}^n [X_i - \widehat{\mathbb{E}}(X)] [X_i - \widehat{\mathbb{E}}(X)]^T.\end{aligned}\tag{30}$$

Now observe that we may write:

$$\begin{aligned}& (n-1)\widehat{\Sigma} \\ &= \sum_{i=1}^n [X_i - \mathbb{E}(X) - (\widehat{\mathbb{E}}(X) - \mathbb{E}(X))] [X_i - \mathbb{E}(X) - (\widehat{\mathbb{E}}(X) - \mathbb{E}(X))]^T \\ &= \sum_{i=1}^n (X_i - \mathbb{E}(X))(X_i - \mathbb{E}(X))^T + n(\widehat{\mathbb{E}}(X) - \mathbb{E}(X))(\widehat{\mathbb{E}}(X) - \mathbb{E}(X))^T \\ &\quad - 2(\widehat{\mathbb{E}}(X) - \mathbb{E}(X)) \sum_{i=1}^n (X_i - \mathbb{E}(X))^T \\ &= n\ddot{\Sigma} - n(\widehat{\mathbb{E}}(X) - \mathbb{E}(X))(\widehat{\mathbb{E}}(X) - \mathbb{E}(X))^T\end{aligned}\tag{31}$$

It follows that:

$$\begin{aligned}& \sqrt{n}(\widehat{\Sigma} - \Sigma) \\ &= \sqrt{n}\left(\frac{n-1}{n-1}\widehat{\Sigma} - \Sigma\right) \\ &= \sqrt{n}\left(\frac{n}{n-1}\ddot{\Sigma} - \frac{n}{n-1}(\widehat{\mathbb{E}}(X) - \mathbb{E}(X))(\widehat{\mathbb{E}}(X) - \mathbb{E}(X))^T - \Sigma\right) \\ &= \frac{n\sqrt{n}}{n-1}\ddot{\Sigma} - \frac{n\sqrt{n}}{n-1}(\widehat{\mathbb{E}}(X) - \mathbb{E}(X))(\widehat{\mathbb{E}}(X) - \mathbb{E}(X))^T - \sqrt{n}\Sigma \\ &= \frac{n\sqrt{n}}{n-1}\ddot{\Sigma} - \frac{n\sqrt{n}}{n-1}(\widehat{\mathbb{E}}(X) - \mathbb{E}(X))(\widehat{\mathbb{E}}(X) - \mathbb{E}(X))^T - \frac{n-1}{n-1}\sqrt{n}\Sigma \\ &= \frac{n\sqrt{n}}{n-1}(\ddot{\Sigma} - \Sigma) - \frac{n\sqrt{n}}{n-1}(\widehat{\mathbb{E}}(X) - \mathbb{E}(X))(\widehat{\mathbb{E}}(X) - \mathbb{E}(X))^T + \frac{\sqrt{n}}{n-1}\Sigma\end{aligned}\tag{32}$$

We are now ready to state the result:

Lemma 1. *Let X_1, \dots, X_n refer to a sequence of i.i.d. random k -vectors. Denote the expectation vector and covariance matrix of X_1 as μ_1 and Σ_1 , respectively. Assume that $\check{\Sigma}_1 = \text{Cov}[v_u((X_1 - \mu_1)(X_1 - \mu_1)^T)]$ is positive definite, where $v_u(M)$ denotes the vectorization of the upper triangular portion of a real symmetric matrix M . Then, we have:*

$$\sqrt{n}(v_u(\widehat{\Sigma}) - v_u(\Sigma_1)) \xrightarrow{d} \mathcal{N}(0, \check{\Sigma}_1).\tag{33}$$

Proof. Consider the quantity $a^T [\sqrt{n}(v_u(\widehat{\Sigma}) - v_u(\Sigma_1))] = \sqrt{n}(a^T v_u(\widehat{\Sigma}) - a^T v_u(\Sigma_1))$ where $a \in \mathbb{R}^{k(k+1)/2} \setminus \{0\}$. Note that $a^T v_u[(X_1 - \mu_1)(X_1 - \mu_1)^T], \dots, a^T v_u[(X_n - \mu_1)(X_n - \mu_1)^T]$ is a sequence of i.i.d. random variables with expectation $a^T v_u(\Sigma_1)$ and variance $a^T \check{\Sigma}_1 a$. Moreover observe that $\check{\Sigma}_1 < \infty$ because $\check{\Sigma}_1$ is positive definite. We can therefore apply the univariate central limit theorem to conclude that:

$$\sqrt{n}(a^T v_u(\check{\Sigma}_1) - a^T v_u(\Sigma_1)) \xrightarrow{d} \mathcal{N}(0, a^T \check{\Sigma}_1 a), \quad (34)$$

where $\check{\Sigma}_1 = \frac{1}{n} \sum_{i=1}^n (X_i - \mu_1)(X_i - \mu_1)^T$. We would however like to claim that:

$$\sqrt{n}(a^T v_u(\widehat{\Sigma}) - a^T v_u(\Sigma_1)) \xrightarrow{d} \mathcal{N}(0, a^T \check{\Sigma}_1 a). \quad (35)$$

In order to prove this, we use 32 and set:

$$\sqrt{n}(a^T v_u(\widehat{\Sigma}) - a^T v_u(\Sigma_1)) = a^T A_n + a^T B_n, \quad (36)$$

where we have:

$$\begin{aligned} A_n &= \frac{n}{n-1} \sqrt{n}(v_u(\check{\Sigma}_1) - v_u(\Sigma)), \\ B_n &= \frac{\sqrt{n}}{n-1} v_u(\Sigma_1) - \frac{n\sqrt{n}}{n-1} v_u[(\widehat{\mathbb{E}}(X) - \mu_1)(\widehat{\mathbb{E}}(X) - \mu_1)^T]. \end{aligned} \quad (37)$$

We already know from 34 that:

$$\sqrt{n}(a^T v_u(\check{\Sigma}_1) - a^T v_u(\Sigma_1)) \xrightarrow{d} \mathcal{N}(0, a^T \check{\Sigma}_1 a). \quad (38)$$

Therefore, so does $a^T A_n$ by Slutsky's lemma, when we view the sequence of constants $\frac{n}{n-1}$ as a sequence of random variables. For $a^T B_n$, we know that:

$$\sqrt{n}(a^T \widehat{\mathbb{E}}(X) - a^T \mu_1) \xrightarrow{d} \mathcal{N}(0, a^T \Sigma_1 a), \quad (39)$$

by viewing $a^T X_1, \dots, a^T X_n$ as a sequence of random variables, noting that $E(X_1 X_1^T) < \infty$ because $\check{\Sigma}_1$ is positive definite and then applying the univariate central limit theorem. We thus have $a^T B_n \xrightarrow{p} 0$. We may then invoke Slutsky's lemma again for $a^T A_n + a^T B_n$ and claim that:

$$\sqrt{n}(a^T v_u(\widehat{\Sigma}) - a^T v_u(\Sigma_1)) \xrightarrow{d} \mathcal{N}(0, a^T \check{\Sigma}_1 a). \quad (40)$$

We conclude the lemma by invoking the Cramer-Wold device. □