

Measuring Dependence Powerfully and Equitably

Yakir A. Reshef*†

YAKIR@SEAS.HARVARD.EDU

*School of Engineering and Applied Sciences
Harvard University
Cambridge, MA 02138, USA*

David N. Reshef*

DNRESHEF@MIT.EDU

*Department of Electrical Engineering and Computer Science
Massachusetts Institute of Technology
Cambridge, MA 02139, USA*

Hilary K. Finucane

HILARYF@MIT.EDU

*Department of Mathematics
Massachusetts Institute of Technology
Cambridge, MA 02139, USA.*

Pardis C. Sabeti**

PSABETI@OEB.HARVARD.EDU

*Department of Organismic and Evolutionary Biology
Harvard University
Cambridge, MA 02138, USA*

Michael Mitzenmacher**

MICHAELM@EECS.HARVARD.EDU

*School of Engineering and Applied Sciences
Harvard University
Cambridge, MA 02138, USA*

* *Co-first author.*

† *To whom correspondence should be addressed.*

** *Co-last author.*

Editor: Edo Airoldi

Abstract

Given a high-dimensional data set, we often wish to find the strongest relationships within it. A common strategy is to evaluate a measure of dependence on every variable pair and retain the highest-scoring pairs for follow-up. This strategy works well if the statistic used (a) has good power to detect non-trivial relationships, and (b) is *equitable*, meaning that for some measure of noise it assigns similar scores to equally noisy relationships regardless of relationship type (e.g., linear, exponential, periodic). In this paper, we define and theoretically characterize two new statistics that together yield an efficient approach for obtaining both power and equitability. To do this, we first introduce a new population measure of dependence and show three equivalent ways that it can be viewed, including as a canonical “smoothing” of mutual information. We then introduce an efficiently computable consistent estimator of our population measure of dependence, and we empirically establish its equitability on a large class of noisy functional relationships. This new statistic has better bias/variance properties and better runtime complexity than a previous heuristic approach. Next, we derive a second, related statistic whose computation is a trivial side-

product of our algorithm and whose goal is powerful independence testing rather than equitability. We prove that this statistic yields a consistent independence test and show in simulations that the test has good power against independence. Taken together, our results suggest that these two statistics are a valuable pair of tools for exploratory data analysis.

Keywords: maximal information coefficient, total information coefficient, equitability, statistical power, mutual information

1. Introduction

The growing dimensionality of today’s data sets has popularized the idea of *hypothesis-generating science*, whereby a data set is used not to test existing hypotheses but rather to help a researcher formulate new ones. A common approach among practitioners is to evaluate some statistic on many candidate variable pairs in a data set, sort the variable pairs from highest-scoring to lowest, and manually examine all the pairs above a threshold score (Storey and Tibshirani, 2003; Emilsson et al., 2008).

A popular class of statistics used for such analyses is *measures of dependence*, i.e., statistics whose population value is zero in cases of statistical independence and non-zero otherwise. Measures of dependence are attractive because they guarantee that asymptotically no non-trivial relationship will erroneously be declared trivial. In the setting of continuous-valued data, which is our focus, there is a long line of fruitful research on such statistics including, e.g., Hoeffding (1948); Rényi (1959); Breiman and Friedman (1985); Paninski (2003); Székely et al. (2007); Gretton et al. (2005); Reshef et al. (2011); Gretton et al. (2012); Lopez-Paz et al. (2013); Heller et al. (2013); Jiang et al. (2015); Heller et al. (2016).

One way to measure the utility of a measure of dependence $\hat{\varphi}$ is *power against independence*, i.e., the power of independence testing based on $\hat{\varphi}$ to detect various types of non-trivial relationships. This is an important goal for data sets that have very few non-trivial relationships, or only very weak relationships that are difficult to detect. Often, however, the number of relationships declared statistically significant by a measure of dependence greatly exceeds the number of relationships that can then be explored further. For example, biological data sets often contain many non-trivial relationships, but further corroborating any one of them may take extensive manual lab work or a study on human or animal subjects. In this case, it is tempting to restrict follow-up to a few relationships with the highest values of $\hat{\varphi}$, but this can skew the direction of follow-up work: if $\hat{\varphi}$ systematically assigns higher scores to, say, linear relationships than to non-linear ones, relatively noisy linear relationships might crowd out strong non-linear relationships from the top-scoring set.

Motivated by this problem, we previously introduced a second way of assessing a measure of dependence, called *equitability* (Reshef et al., 2011). Informally, an equitable statistic is one that, for some measure of relationship strength, assigns similar scores to equally strong relationships regardless of relationship type. For instance, we may want our measure of dependence to also have the property that on noisy functional relationships it assigns similar scores to relationships with the same R^2 , i.e., the squared Pearson correlation between the observed y-values and the x-values passed through the underlying function in question (Reshef et al., 2011). Or, alternatively, we may want the value of our statistic to tell us about the proportion of points coming from the deterministic component of a mixture containing part signal and part uniform noise (Ding and Li, 2013). Defining measures of dependence that achieve good equitability with respect to interesting measures of relationship strength

is a new and challenging problem, with a number of different formalizations. (See, e.g., [Reshef et al., 2015b](#) and [Ding and Li, 2013](#) cited above, as well as [Kinney and Atwal, 2014](#) along with associated technical comments [Reshef et al., 2014](#) and [Murrell et al., 2014](#).) A companion paper to this work ([Reshef et al., 2015b](#)) presents a general formalization that unifies these.

In this paper, we introduce and theoretically characterize two new measures of dependence that we empirically show to have good equitability with respect to R^2 and power against independence, respectively. We begin by introducing a new population measure of dependence called MIC_* . Given a pair of jointly distributed random variables (X, Y) , $\text{MIC}_*(X, Y)$ is the supremum, over all finite grids G imposed on the support of (X, Y) , of the mutual information of the discrete distribution induced by (X, Y) on the cells of G , subject to a regularization based on the resolution of G . We prove three results, each of which gives a different way that this population quantity can be viewed.

1. MIC_* is the population value of the maximal information coefficient (MIC), a statistic introduced in [Reshef et al. \(2011\)](#) that is empirically highly equitable with respect to R^2 on a large class of noisy functional relationships. Simple corollaries of this result simplify and strengthen many of the theoretical results proven in [Reshef et al. \(2011\)](#) about MIC.
2. MIC_* is a minimal smoothing of mutual information, in the sense that the regularization in the definition of MIC_* renders it uniformly continuous as a function of random variables with respect to statistical distance, and no “smaller” regularization achieves continuity. This result yields as a corollary that mutual information by itself is not continuous with respect to statistical distance.
3. MIC_* is the supremum of an infinite sequence defined in terms of optimal (one-dimensional) partitions of the marginal distributions of (X, Y) rather than optimal (two-dimensional) grids imposed on the joint distribution. This characterization greatly simplifies computation.

After proving these three results, we leverage them to introduce efficient algorithms both for approximating MIC_* in practice and for estimating it consistently from a finite sample. We first provide an efficient algorithm that in many cases allows for computation to arbitrary precision of the MIC_* of a pair of random variables whose joint density is known. We then introduce a statistic, called MIC_e , that we prove is a consistent estimator of MIC_* . In contrast to the MIC statistic from [Reshef et al. \(2011\)](#), for which no efficient algorithm is known and a heuristic algorithm is used in practice, MIC_e is efficiently computable. It has a better runtime complexity than the heuristic algorithm currently in use for computing the original MIC statistic, and is orders of magnitude faster in practice.

With a consistent and fast estimator for MIC_* in hand, we turn to empirical analysis of its performance. Specifically, we show through simulation that MIC_e has better bias/variance properties than the heuristic algorithm used in [Reshef et al. \(2011\)](#) for computing MIC, which has no theoretical convergence guarantees. Our analysis also reveals that the main parameter of MIC_e can be used to tune statistical performance toward either stronger or weaker relationships in general. After studying the bias/variance properties of MIC_e , we

then demonstrate via simulation that it outperforms currently available methods in terms of equitability with respect to R^2 on a broad set of noisy functional relationships. We show this performance advantage both on the set of functional relationships analyzed in [Reshef et al. \(2011\)](#) as well as on a large set of randomly chosen noisy functional relationships.

We choose in this paper to analyze equitability specifically with respect to R^2 , rather than some other notion of relationship strength, because R^2 on noisy functional relationships is a simple measure with broad familiarity and intuitive interpretation among practitioners. Of course, it is also important to develop measures of dependence that are equitable with respect to notions of relationship strength besides R^2 or on families of relationships besides noisy functional relationships; however, our focus here remains on the “simple” case of R^2 on noisy functional relationships.

Importantly, we note that although there are methods for directly estimating the R^2 of a noisy functional relationship via nonparametric regression (see, e.g., [Cleveland and Devlin, 1988](#); [Stone, 1977](#)), those methods are not applicable in the context of equitability because they are not measures of dependence. That is, because non-parametric regression methods *assume* a functional form for the relationship in question, they can give trivial scores to non-functional relationships, even in the large-sample limit. (A simple example of this is a uniform distribution over a circle, whose regression function is constant.) In contrast, a *measure of dependence* is guaranteed never to make this “mistake”. A measure of dependence that is equitable with respect to R^2 can therefore be viewed either as an “upgraded” measure of dependence that also comes with some of the interpretability properties of non-parametric regression, or as an “upgraded” approximate non-parametric regression method that also has the robustness properties of a measure of dependence.

The main strength of MIC_e is equitability rather than power to reject a null hypothesis of independence. In some settings, though, it may be more important to focus on good power against independence. We therefore introduce here a statistic closely related to MIC_e called the total information coefficient and denoted TIC_e . We prove the consistency of testing for independence using TIC_e , and show via simulations that it achieves excellent power in practice, performing comparably to or better than current methods on an index suite of relationships from [Simon and Tibshirani \(2012\)](#). Because TIC_e arises naturally as a side-product of the computation of MIC_e , it is available “for free” once MIC_e has been computed. This leads us to propose a data analysis strategy consisting of first using TIC_e to filter out non-significant relationships, and then ranking the remaining ones using the simultaneously computed values of MIC_e .

In addition to the companion paper [Reshef et al. \(2015b\)](#), which focuses on the theory behind equitability, this paper is accompanied by a second companion work ([Reshef et al., 2015a](#)) that explores in detail the empirical performance of the methods introduced here. That paper compares MIC_e and TIC_e to several leading measures of dependence ([Kraskov et al., 2004](#); [Székely and Rizzo, 2009](#); [Heller et al., 2013, 2016](#); [Gretton et al., 2005](#); [Breiman and Friedman, 1985](#); [Lopez-Paz et al., 2013](#)) on a broad range of relationship types under many different sampling and noise models, finding that the equitability with respect to R^2 of MIC_e and the power of independence testing using TIC_e are both state-of-the-art on the relationships examined. It also shows that these methods can be computed very fast in practice.

Taken together, our results shed significant light on the theory behind the maximal information coefficient, and suggest that TIC_e and MIC_e are a useful pair of methods for data exploration. Specifically, they point to joint use of these two statistics to filter and then rank relationships as a fast, practical way to explore large data sets by measuring dependence both powerfully and equitably.

2. Preliminaries

We work extensively in this paper with grids and discrete distributions over their cells. Given a grid G and a point (x, y) , we define the function $\text{row}_G(y)$ to be the row of G containing y and we define $\text{col}_G(x)$ analogously. For a pair (X, Y) of jointly distributed random variables, we write $(X, Y)|_G$ to denote $(\text{col}_G(X), \text{row}_G(Y))$, and we use $I((X, Y)|_G)$ to denote the discrete mutual information (Cover and Thomas, 2006; Csiszár and Shields, 2004; Csiszár, 2008) between $\text{col}_G(X)$ and $\text{row}_G(Y)$. Given a finite sample D from the distribution of (X, Y) , we sometimes use D to refer both to the set of points in the sample as well as to a point chosen uniformly at random from D . In the latter case, it will then make sense to talk about, e.g., $D|_G$ and $I(D|_G)$.

For natural numbers k and ℓ , we use $G(k, \ell)$ to denote the set of all k -by- ℓ grids (possibly with empty rows/columns). A grid G is an equipartition of (X, Y) if all the rows of $(X, Y)|_G$ have the same probability mass, and all the columns do as well. We also use the term equipartition in the analogous way for one-dimensional partitions into just rows or columns. For a one-dimensional partition P into rows and a one-dimensional partition Q into columns, we write (P, Q) to refer to the grid constructed from these two partitions. When a partition P can be obtained from a partition P' by addition of separators alone, we write $P' \subset P$.

Finally, let us establish some notation for infinite matrices. We use m^∞ to denote the space of infinite matrices equipped with the supremum norm. Given a matrix $A \in m^\infty$, we often examine only the k, ℓ -th entries of A for which $k\ell \leq i$ for some i . Thus, for $i \in \mathbb{Z}^+$, we define the projection $r_i : m^\infty \rightarrow m^\infty$ via

$$r_i(A)_{k,\ell} = \begin{cases} A_{k,\ell} & k\ell \leq i \\ 0 & k\ell > i \end{cases} .$$

Unless noted otherwise, all logarithms are to base 2.

3. The Population Maximal Information Coefficient MIC_*

In this section, we define and characterize the population maximal information coefficient MIC_* . We begin by defining the population quantity $\text{MIC}_*(X, Y)$ for a pair of jointly distributed random variables (X, Y) . We then show three different ways to characterize this population quantity: first, as the large-sample limit of the statistic MIC from Reshef et al. (2011); second, as a minimally smoothed version of mutual information; and third, as the supremum of an infinite sequence defined in terms of optimal one-dimensional partitions of the marginal distributions of (X, Y) . We conclude the section by showing how the third characterization leads to an efficient approach for approximating MIC_* in practice from the density of (X, Y) .

3.1 Defining MIC_*

The population maximal information coefficient can be defined in several equivalent ways, as we will see later. For now, we begin with the simplest definition.

Definition 1 *Let (X, Y) be jointly distributed random variables. The population maximal information coefficient (MIC_*) of (X, Y) is defined by*

$$MIC_*(X, Y) = \sup_G \frac{I((X, Y)|_G)}{\log \|G\|}$$

where $\|G\|$ denotes the minimum of the number of rows of G and the number of columns of G .

Given that $I(X, Y) = \sup_G I((X, Y)|_G)$ (see, e.g., Chapter 8 of [Cover and Thomas 2006](#)), this can be viewed as a regularized version of mutual information that penalizes complicated grids and ensures that the result falls between zero and one.

Before we continue, we state one simple equivalent definition of MIC_* that is useful for the results in this section. This definition views MIC_* as the supremum of a matrix called the *population characteristic matrix*, defined below.

Definition 2 *Let (X, Y) be jointly distributed random variables. Let*

$$I^*((X, Y), k, \ell) = \max_{G \in \mathcal{G}(k, \ell)} I((X, Y)|_G).$$

The population characteristic matrix of (X, Y) , denoted by $M(X, Y)$, is defined by

$$M(X, Y)_{k, \ell} = \frac{I^*((X, Y), k, \ell)}{\log \min\{k, \ell\}}$$

for $k, \ell > 1$.

It is easy to see the following:

Proposition 3 *Let (X, Y) be jointly distributed random variables. We have*

$$MIC_*(X, Y) = \sup M(X, Y)$$

where $M(X, Y)$ is the population characteristic matrix of (X, Y) .

The population characteristic matrix is so named because just as MIC_* , the supremum of this matrix, captures a sense of relationship strength, other properties of this matrix correspond to different properties of relationships. For instance, later in this paper we introduce an additional property of the characteristic matrix, the total information coefficient, that is useful for testing for the presence or absence of a relationship rather than quantifying relationship strength.

3.2 First Alternate Characterization: MIC_* Is the Population Value of MIC

With MIC_* defined, we now state our first alternate characterization of it, as the large-sample limit of the statistic MIC introduced in Reshef et al. (2011). We begin by first reproducing a description of MIC from Reshef et al. (2011), via the two definitions below.

Definition 4 (Reshef et al., 2011) Let $D \subset \mathbb{R}^2$ be a set of ordered pairs. The sample characteristic matrix $\widehat{M}(D)$ of D is defined by

$$\widehat{M}(D)_{k,\ell} = \frac{I^*(D, k, \ell)}{\log \min\{k, \ell\}}.$$

Definition 5 (Reshef et al., 2011) Let $D \subset \mathbb{R}^2$ be a set of n ordered pairs, and let $B : \mathbb{Z}^+ \rightarrow \mathbb{Z}^+$. We define

$$MIC_B(D) = \max_{k\ell \leq B(n)} \widehat{M}(D)_{k,\ell}.$$

where the function $B(n)$ is specified by the user. In Reshef et al. (2011), it was suggested that $B(n)$ be chosen to be n^α for some constant α in the range of 0.5 to 0.8. (The statistics we introduce later will have an analogous parameter; see Section 4.4.1.)

We show the following result about convergence of functions of the sample characteristic matrix to their population counterparts, a consequence of which is the convergence of MIC to MIC_* . (In the theorem statement below, recall that m^∞ is the space of infinite matrices equipped with the supremum norm, and given a matrix A the projection r_i zeros out all the entries $A_{k,\ell}$ for which $k\ell > i$.)

Theorem 6 Let $f : m^\infty \rightarrow \mathbb{R}$ be uniformly continuous, and assume that $f \circ r_i \rightarrow f$ pointwise. Then for every random variable (X, Y) , we have

$$(f \circ r_{B(n)}) \left(\widehat{M}(D_n) \right) \rightarrow f(M(X, Y))$$

in probability where D_n is a sample of size n from the distribution of (X, Y) , provided $\omega(1) < B(n) \leq O(n^{1-\varepsilon})$ for some $\varepsilon > 0$.

Proof See Appendix A. ■

Since the supremum of a matrix is uniformly continuous as a function on m^∞ and can be realized as the limit of maxima of larger and larger segments of the matrix, this theorem yields our claim about MIC_* as a corollary.

Corollary 7 MIC_B is a consistent estimator of MIC_* provided $\omega(1) < B(n) \leq O(n^{1-\varepsilon})$ for some $\varepsilon > 0$.

Though Theorem 6 is proven in Appendix A, we provide here some intuition for why it should hold as well as a description of the obstacles that must be overcome in the proof.

For concreteness, suppose f is the supremum function. To see why the theorem should hold, fix a random variable (X, Y) and let D be a sample of size n from its distribution. It is known that for a fixed grid G $I(D|_G)$ is a consistent estimator of $I((X, Y)|_G)$ (Roulston,

1999; Paninski, 2003). We might therefore expect $I^*(D, k, \ell)$ to be a consistent estimator of $I^*((X, Y), k, \ell)$ as well. And if $I^*(D, k, \ell)$ is a consistent estimator of $I^*((X, Y), k, \ell)$, then we might expect the maximum of the sample characteristic matrix (which just consists of normalized I^* terms) to be a consistent estimator of the supremum of the true characteristic matrix.

These intuitions turn out to be true, but there are two reasons they are non-trivial to prove. First, consistency for I^* does not follow from abstract considerations since the supremum of an infinite set of estimators is not necessarily a consistent estimator of the supremum of the estimands.¹ Second, consistency of I^* alone does not suffice to show that the maximum of the sample characteristic matrix converges to MIC_* . In particular, if $B(n)$ grows too quickly, and the convergence of $I^*(D, k, \ell)$ to $I^*((X, Y), k, \ell)$ is slow, inflated values of MIC can result. To see this, notice that if $B(n) = \infty$ then $\text{MIC} = 1$ for uniformly generated noise at any finite sample size, even though each individual entry of the sample characteristic matrix converges to its true value eventually.

The technical heart of the proof is overcoming these obstacles by using the dependencies between the quantities $I(D|_G)$ for different grids G to not only show the consistency of $I^*(D, k, \ell)$ but then to quantify how quickly $I^*(D, k, \ell)$ converges to $I^*((X, Y), k, \ell)$.

3.3 Second Alternate Characterization: MIC_* Is a Minimally Smoothed Mutual Information

We now describe a second equivalent view of MIC_* . Recall that for a pair of jointly distributed random variables (X, Y) , we defined $\text{MIC}_*(X, Y)$ as

$$\text{MIC}_*(X, Y) = \sup_G \frac{I((X, Y)|_G)}{\log \|G\|}$$

where $\|G\|$ denotes the minimum of the number of rows of G and the number of columns of G . As we discussed in Section 3.1, the mutual information $I(X, Y)$ is also a supremum, namely

$$I(X, Y) = \sup_G I((X, Y)|_G).$$

and so MIC_* can be viewed as a regularized version of I . It is natural to ask whether the regularization in the definition of MIC_* has any smoothing effect on I . In this sub-section we show first that it does, in the sense that MIC_* is uniformly continuous as a function of random variables with respect to the metric of statistical distance,² and second that the regularization by $\log \|G\|$ is in some sense the minimal one necessary for achieving any sort of continuity. As a corollary, we obtain that I by itself is not continuous as a function of

-
1. If $\hat{\theta}_1, \dots, \hat{\theta}_k$ is a finite set of estimators, then a union bound shows that the random variable $(\hat{\theta}_1(D), \dots, \hat{\theta}_k(D))$ converges in probability to $(\theta_1, \dots, \theta_k)$ with respect to the supremum metric. The continuous mapping theorem then gives the desired result. However, if the set of estimators is infinite, the union bound cannot be employed. And indeed, if we let $\theta_1 = \dots = \theta_k = 0$, and let $\hat{\theta}_i(D_n) = i/n$ deterministically, then each $\hat{\theta}_i$ is a consistent estimator of θ_i , but since the set $\{\hat{\theta}_1(D_n), \hat{\theta}_2(D_n), \dots\} = \{1/n, 2/n, \dots\}$ is unbounded, $\sup_i \hat{\theta}_i(D_n) = \infty$ for every n .
 2. Recall that the statistical distance between random variables A and B is defined as $\sup_T |\mathbf{P}(A \in T) - \mathbf{P}(B \in T)|$. When A and B have probability density functions or probability mass functions, this equals one-half of the L^1 distance between those functions.

random variables with respect to the metric of statistical distance. This provides a view of MIC_* as a canonical smoothing of I that yields continuity.

Formally, let $\mathcal{P}(\mathbb{R}^2)$ denote the space of random variables supported on \mathbb{R}^2 equipped with the metric of statistical distance. Our first claim is that as a function defined on $\mathcal{P}(\mathbb{R}^2)$, MIC_* is uniformly continuous. We prove this claim by establishing a stronger result: the uniform continuity of the characteristic matrix $M(X, Y)$. Specifically, by showing that the family of maps corresponding to each individual entry of the characteristic matrix is uniformly equicontinuous, we obtain the following result.

Theorem 8 *The map from $\mathcal{P}(\mathbb{R}^2)$ to m^∞ defined by $(X, Y) \mapsto M(X, Y)$ is uniformly continuous.*

Proof See Appendix B. ■

Since the supremum is a uniformly continuous function on m^∞ , Theorem 8 yields the following corollary.

Corollary 9 *The map $(X, Y) \mapsto \text{MIC}_*(X, Y)$ is uniformly continuous.*

Similar corollaries exist for any uniformly continuous function of the characteristic matrix.

Interestingly, Theorem 8 relies crucially on the normalization in the definition of the characteristic matrix. This is not a coincidence: as the following proposition shows, any normalization that is meaningfully smaller than the one in the definition of the characteristic matrix will cause the matrix to contain a discontinuity as a function on $\mathcal{P}(\mathbb{R}^2)$.

Proposition 10 *For some function $N(k, \ell)$, let M^N be the characteristic matrix with normalization N , i.e.,*

$$M^N(X, Y)_{k, \ell} = \frac{I^*((X, Y), k, \ell)}{N(k, \ell)}.$$

If $N(k, \ell) = o(\log \min\{k, \ell\})$ along some infinite path in $\mathbb{N} \times \mathbb{N}$, then M^N and $\sup M^N$ are not continuous as functions of $\mathcal{P}([0, 1] \times [0, 1]) \subset \mathcal{P}(\mathbb{R}^2)$.

Proof See Appendix C. ■

The above proposition implies that the “smoothing” that MIC_* applies to mutual information is necessary in some sense. In particular, one corollary of the proposition is that mutual information with no smoothing will contain a discontinuity.

Corollary 11 *Mutual information is not continuous on $\mathcal{P}([0, 1] \times [0, 1]) \subset \mathcal{P}(\mathbb{R}^2)$.*

Proof Mutual information is the supremum of M^N with $N \equiv 1$. ■

The same result can also be shown for the squared Linfoot correlation (Speed, 2011; Linfoot, 1957), which equals $1 - 2^{-2I}$ where I represents mutual information. Thus, though the Linfoot correlation smoothes the mutual information enough to cause it to lie in the unit interval, it does not smooth the mutual information sufficiently to cause it to be continuous.

As we remarked previously, these results, when contrasted with the uniform continuity of MIC_* , allow us to view the latter as a canonical “minimally smoothed” version of mutual information that is uniformly continuous. This view gives a meaningful interpretation to the normalization used in MIC_* . Understanding MIC_* as having smoothness properties not shared by mutual information also suggests that estimators of MIC_* may have better statistical properties than estimators of ordinary mutual information. This is consistent with a recent hardness-of-estimation result for mutual information in [Ding and Li \(2013\)](#) and is also borne out empirically in [Reshef et al. \(2015a\)](#).

3.4 Third Alternate Characterization: MIC_* Is the Supremum of the Boundary of the Characteristic Matrix

We now show the third alternate view of MIC_* : that it can be equivalently defined as the supremum over a *boundary* of the characteristic matrix rather than as a supremum over all of the entries of the matrix. This characterization of MIC_* will serve as the foundation both for our approach to approximating $\text{MIC}_*(X, Y)$ as well as the new estimator of MIC_* that we introduce later in this paper.

We begin by defining what we mean by the boundary of the characteristic matrix. Our definition rests on the following observation.

Proposition 12 *Let M be a population characteristic matrix. Then for $\ell \geq k$, $M_{k,\ell} \leq M_{k,\ell+1}$.*

Proof Let (X, Y) be the random variable in question. Since we can always let a row/column be empty, we know that $I^*((X, Y), k, \ell) \leq I^*((X, Y), k, \ell + 1)$. And since $\ell, \ell + 1 \geq k$, we know that $M_{k,\ell} = I^*((X, Y), k, \ell) / \log k \leq I^*((X, Y), k, \ell + 1) / \log k = M_{k,\ell+1}$. ■

Since the entries of the characteristic matrix are bounded, the monotone convergence theorem then gives the following corollary. In the corollary and henceforth, we let $M_{k,\uparrow} = \lim_{\ell \rightarrow \infty} M_{k,\ell}$ and define $M_{\uparrow,\ell}$ similarly.

Corollary 13 *Let M be a population characteristic matrix. Then $M_{k,\uparrow}$ exists, is finite, and equals $\sup_{\ell \geq k} M_{k,\ell}$. The same is true for $M_{\uparrow,\ell}$.*

The above corollary allows us to define the *boundary* of the characteristic matrix.

Definition 14 *Let M be a population characteristic matrix. The boundary of M is the set*

$$\partial M = \{M_{k,\uparrow} : 1 < k < \infty\} \cup \{M_{\uparrow,\ell} : 1 < \ell < \infty\}.$$

The theorem below then gives a relationship between the boundary of the characteristic matrix and MIC_* .

Theorem 15 *Let (X, Y) be a random variable. We have*

$$\text{MIC}_*(X, Y) = \sup \partial M(X, Y)$$

where $M(X, Y)$ is the population characteristic matrix of (X, Y) .

Proof The following argument shows that every entry of M is at most $\sup \partial M$: fix a pair (k, ℓ) and notice that either $k \leq \ell$, in which case $M_{k,\ell} \leq M_{k,\uparrow}$, or $\ell \leq k$, in which case $M_{k,\ell} \leq M_{\uparrow,\ell}$. Thus, $\text{MIC}_* \leq \sup\{M_{\uparrow,\ell}\} \cup \{M_{k,\uparrow}\} = \sup \partial M$.

On the other hand, Corollary 13 shows that each element of ∂M is a supremum over some elements of M . Therefore, $\sup \partial M$, being a supremum over suprema of elements of M , cannot exceed $\sup M = \text{MIC}_*$. ■

3.5 Approximating MIC_* in Practice

The importance of the characterization in Theorem 15 from the previous sub-section is computational. Specifically, elements of the boundary of the characteristic matrix can be expressed in terms of a maximization over (one-dimensional) partitions rather than (two-dimensional) grids, the former being much quicker to compute exactly. This is stated in the theorem below.

Theorem 16 *Let M be a population characteristic matrix. Then $M_{k,\uparrow}$ equals*

$$\max_{P \in \mathcal{P}(k)} \frac{I(X, Y|_P)}{\log k}$$

where $\mathcal{P}(k)$ denotes the set of all partitions of size at most k .

Proof See Appendix D. ■

To formally state how this will help us from an algorithmic standpoint, we note that Theorems 15 and 16 above together give the following corollary.

Corollary 17 *Let (X, Y) be a random variable, and let \mathbb{P} be the set of finite-size partitions. Then*

$$\text{MIC}_*(X, Y) = \sup \left\{ \frac{I(X, Y|_P)}{\log |P|} : P \in \mathbb{P} \right\} \cup \left\{ \frac{I(X|_P, Y)}{\log |P|} : P \in \mathbb{P} \right\}$$

where $|P|$ is the number of bins in the partition P .

We can exploit the fact that the expressions in the above corollary involve maximization only over one-dimensional partitions rather than two-dimensional grids to give an algorithm for computing elements of the boundary of the characteristic matrix to arbitrary precision, and by extension an approach to approximating MIC_* in practice. To do so, we utilize as a subroutine a dynamic programming algorithm from Reshef et al. (2011) called OPTIMIZEX-AXIS. Before continuing, we therefore give a brief overview of that algorithm.

Overview of OPTIMIZEXAXIS algorithm from Reshef et al. (2011). The OPTIMIZEX-AXIS algorithm takes as input a set D of n data points, a fixed partition into columns³ Q of size ℓ , a “master” partition into rows Π , and a number k . The algorithm returns, for

3. Despite its name, the OPTIMIZEXAXIS algorithm can be used to optimize a partition of either axis. In our description of the algorithm here, we choose to describe the algorithm as it would work for optimizing a partition of the y -axis rather than the x -axis. This is for notational coherence of this paper only.

$2 \leq i \leq k$, the partition into rows $P_i \subset \Pi$ that maximizes the mutual information of $D|_{(P_i, Q)}$ among all sub-partitions of Π of size at most i . The algorithm works by exploiting the fact that, conditioned on the location y of the top-most line of P_i , the optimization of the rest of P_i can be formulated as a sub-problem that depends only on the data points below y . The algorithm uses dynamic programming to store and reuse solutions to these subproblems, resulting in a runtime of $O(|\Pi|^2 k \ell)$. If a black-box algorithm is used to compute each required mutual information in time at most T , then the runtime of the algorithm can be shown to be $O(Tk|\Pi|)$.

The following theorem shows that the theory developed about the boundary of the characteristic matrix, together with OPTIMIZEXAXIS, yields an efficient algorithm for computing entries of the boundary to arbitrary precision.

Theorem 18 *Given a random variable (X, Y) , $M_{k, \uparrow}$ (resp. $M_{\uparrow, \ell}$) is computable to within an additive error of $O(k\varepsilon \log(1/(k\varepsilon))) + E$ (resp. $O(\ell\varepsilon \log(1/(\ell\varepsilon))) + E$) in time $O(kT(E)/\varepsilon)$ (resp. $O(\ell T(E)/\varepsilon)$), where $T(E)$ is the time required to numerically compute the mutual information of a continuous distribution to within an additive error of E .*

Proof See Appendix E. ■

The algorithm proposed in Theorem 18 gives us a polynomial-time method for computing any finite subset of the boundary ∂M of the population characteristic matrix $M(X, Y)$ of a random variable (X, Y) . Thus, if we have some k_0, ℓ_0 such that the maximum of the finite subset $\{M_{k, \uparrow}, M_{\uparrow, \ell} : k \leq k_0, \ell \leq \ell_0\}$ of ∂M will be ε -close to the supremum of the entire set ∂M , we can compute $\text{MIC}_*(X, Y)$ to within an error of ε . Though we usually do not have precise knowledge of k_0 and ℓ_0 , for many distributions it is often easy to make very conservative educated guesses for them, in which case this algorithm allows us to approximate $\text{MIC}_*(X, Y)$ very well in practice.

Being able to compute $\text{MIC}_*(X, Y)$ to arbitrary precision in some cases has two main advantages. The first advantage is that it allows us to assess in simulations the large-sample properties of MIC_* independent of any estimator. This is done in the companion paper (Reshef et al., 2015a), which shows that MIC_* achieves high equitability with respect to R^2 on a set of noisy functional relationships thereby confirming that statistically efficient estimation of MIC_* is a worthwhile goal.

The second advantage is that we can empirically assess the bias, variance, and expected squared error of estimators of MIC_* by taking a distribution, computing MIC_* , and then comparing the result to estimates of it based on finite samples. In the next section, we introduce a new estimator MIC_e of MIC_* and carry out such an analysis to compare its statistical properties to those of the statistic MIC from Reshef et al. (2011).

4. Estimating MIC_* with MIC_e

As we have shown, MIC_* is the population value of the statistic MIC introduced in Reshef et al. (2011). However, though consistent, the statistic MIC is not known to be efficiently computable and in Reshef et al. (2011) a heuristic approximation algorithm called APPROX-MIC was computed instead. In this section, we leverage the theory we have developed here

to introduce a new estimator of MIC_* that is both consistent and efficiently computable. The new estimator, called MIC_e , has better runtime complexity even than the heuristic APPROX-MIC algorithm, and runs orders of magnitude faster in practice.

The estimator MIC_e is based on one of the alternate characterizations of MIC_* proven in the previous section. Namely, if MIC_* can be viewed as the supremum of the *boundary* of the characteristic matrix rather than of the entire matrix, then only the boundary of the matrix must be accurately estimated in order to estimate MIC_* . This has the advantage that, whereas computing individual entries of the sample characteristic matrix involves finding optimal (two-dimensional) grids, estimating entries of the boundary requires us only to find optimal (one-dimensional) partitions. While the former problem is computationally difficult, the latter can be solved using the dynamic programming algorithm from [Reshef et al. \(2011\)](#) that we also employed in Section 3.5 to compute MIC_* to arbitrary precision in the large-sample limit.

We formalize this idea via a new object called the *equicharacteristic matrix*, which we denote by $[M]$. The difference between $[M]$ and the characteristic matrix M is as follows: while the k, ℓ -th entry of M is computed from the maximal achievable mutual information using any k -by- ℓ grid, the k, ℓ -th entry of $[M]$ is computed from the maximal achievable mutual information using any k -by- ℓ grid that equipartitions the dimension with more rows/columns. (See Figure 1.) Despite this difference, as the equipartition in question gets finer and finer it becomes indistinguishable from an optimal partition of the same size. This intuition can be formalized to show that the boundary of $[M]$ equals the boundary of M , and therefore that $\sup[M] = \sup M = \text{MIC}_*$. It will then follow that estimating $[M]$ and taking the supremum—as we did with M in the case of MIC —yields a consistent estimate of MIC_* .

4.1 The Equicharacteristic Matrix

We now define the equicharacteristic matrix and show that its supremum is indeed MIC_* . To do so, we first define a version of I^* that equipartitions the dimension with more rows/columns. Note that in the definition, brackets are used to indicate the presence of an equipartition.

Definition 19 *Let (X, Y) be jointly distributed random variables. Define*

$$I^*((X, Y), k, [\ell]) = \max_{G \in G(k, [\ell])} I((X, Y)|_G)$$

where $G(k, [\ell])$ is the set of k -by- ℓ grids whose y -axis partition is an equipartition of size ℓ . Define $I^*((X, Y), [k], \ell)$ analogously.

Define $I^{[*]}((X, Y), k, \ell)$ to equal $I^*((X, Y), k, [\ell])$ if $k < \ell$ and $I^*((X, Y), [k], \ell)$ otherwise.

We now define the equicharacteristic matrix in terms of $I^{[*]}$. In the definition below, we continue our convention of using brackets to denote the presence of equipartitions.

Definition 20 *Let (X, Y) be jointly distributed random variables. The population equicharacteristic matrix of (X, Y) , denoted by $[M](X, Y)$, is defined by*

$$[M](X, Y)_{k, \ell} = \frac{I^{[*]}((X, Y), k, \ell)}{\log \min\{k, \ell\}}$$

for $k, \ell > 1$.

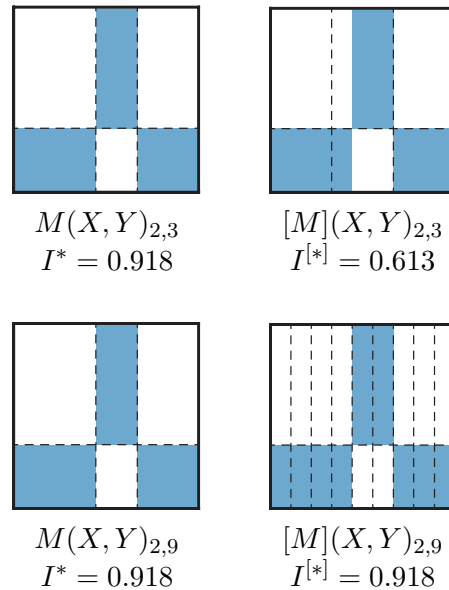


Figure 1: A schematic illustrating the difference between the characteristic matrix M and the equicharacteristic matrix $[M]$. (Top) When restricted to 2 rows and 3 columns, the characteristic matrix M is computed from the optimal 2-by-3 grid. In contrast, the equicharacteristic matrix $[M]$ still optimizes the smaller partition of size 2 but is restricted to have the larger partition be an equipartition of size 3. This results in a lower mutual information of 0.613. (Bottom) When 9 columns are allowed instead of 3, the grid found by the characteristic matrix does not change, since the grid with 3 columns was already optimal. However, now the equicharacteristic matrix uses an equipartition into columns of size 9, whose resolution is able to fully capture the dependence between X and Y .

The boundary of the equicharacteristic matrix can be defined via a limit in the same way as the characteristic matrix. We then have the following theorem.

Theorem 21 *Let (X, Y) be jointly distributed random variables. Then $\partial[M] = \partial M$.*

Proof See Appendix F. ■

Since every entry of the equicharacteristic matrix is dominated by some entry on its boundary, the equivalence of $\partial[M]$ and ∂M yields the following corollary as a simple consequence.

Corollary 22 *Let (X, Y) be jointly distributed random variables. Then $\sup[M](X, Y) = MIC_*(X, Y)$.*

4.2 The Estimator MIC_e

With the equicharacteristic matrix defined, we can now define our new estimator MIC_e in terms of the sample equicharacteristic matrix, analogously to the way we defined MIC in terms of the sample characteristic matrix.

Definition 23 Let $D \subset \mathbb{R}^2$ be a set of ordered pairs. The sample equicharacteristic matrix $[\widehat{M}](D)$ of D is defined by

$$[\widehat{M}](D)_{k,\ell} = \frac{I^{[*]}(D, k, \ell)}{\log \min\{k, \ell\}}.$$

Definition 24 Let $D \subset \mathbb{R}^2$ be a set of n ordered pairs, and let $B : \mathbb{Z}^+ \rightarrow \mathbb{Z}^+$. We define

$$\text{MIC}_{e,B}(D) = \max_{k\ell \leq B(n)} [\widehat{M}](D)_{k,\ell}.$$

With the equivalence between the boundary of the characteristic matrix and that of the equicharacteristic matrix established, it is straightforward to show that MIC_e is a consistent estimator of MIC_* via arguments similar to those we applied in the case of MIC . (See Appendix G.) Specifically, we show the following theorem, an analogue of Theorem 6.

Theorem 25 Let $f : m^\infty \rightarrow \mathbb{R}$ be uniformly continuous, and assume that $f \circ r_i \rightarrow f$ pointwise. Then for every random variable (X, Y) , we have

$$(f \circ r_{B(n)}) \left([\widehat{M}](D_n) \right) \rightarrow f([M](X, Y))$$

in probability where D_n is a sample of size n from the distribution of (X, Y) , provided $\omega(1) < B(n) \leq O(n^{1-\varepsilon})$ for some $\varepsilon > 0$.

By setting $f([M]) = \sup[M]$, we then obtain as a corollary the consistency of MIC_e .

Corollary 26 $\text{MIC}_{e,B}$ is a consistent estimator of MIC_* provided $\omega(1) < B(n) \leq O(n^{1-\varepsilon})$ for some $\varepsilon > 0$.

As with the statistic MIC , the statistic MIC_e requires the user to specify a function $B(n)$ to use. While the theory suggests that any function of the form $B(n) = n^\alpha$ suffices provided $0 < \alpha < 1$, different values of α may yield different finite-sample properties. We study the empirical performance of MIC_e for different choices of $B(n)$ in Section 4.4 and point the reader to specific recommendations for practical use in Section 4.4.1.

4.3 Computing MIC_e

Both MIC and MIC_e are consistent estimators of MIC_* . The difference between them is that while MIC can currently be computed efficiently only via a heuristic approximation, MIC_e can be computed exactly, very efficiently, via an approach similar to the one used for approximating MIC_* involving the `OPTIMIZEXAXIS` subroutine. We now describe the details of this approach.

Recall that, given a fixed x-axis partition Q into ℓ columns, a set of n data points, a “master” y-axis partition Π , and a number k , the `OPTIMIZEXAXIS` subroutine finds, for

every $2 \leq i \leq k$, a y-axis partition $P_i \subset \Pi$ of size at most i that maximizes the mutual information induced by the grid (P_i, Q) . The algorithm does this in time $O(|\Pi|^2 k \ell)$. (For more discussion of OPTIMIZEXAXIS, see Section 3.5)

In the pair of theorems below, we show two ways that OPTIMIZEXAXIS can be used to compute MIC_e efficiently. In the proofs of both theorems, we neglect issues of divisibility, e.g., we often write $B/2$ rather than $\lfloor B/2 \rfloor$. This does not affect the results.

Theorem 27 *There exists an algorithm EQUICHAR that, given a sample D of size n and some $B \in \mathbb{Z}^+$, computes the portion $r_{B(n)}(\widehat{M}(D))$ of the sample equicharacteristic matrix in time $O(n^2 B^2)$, which equals $O(n^{4-2\varepsilon})$ for $B(n) = O(n^{1-\varepsilon})$ with $\varepsilon > 0$.*

Proof We describe the algorithm and simultaneously bound its runtime. We do so only for the k, ℓ -th entries of $\widehat{M}(D)$ satisfying $k \leq \ell, k\ell \leq B$. This suffices, since by symmetry computing the rest of the required entries at most doubles the runtime.

To compute $\widehat{M}(D)_{k,\ell}$ with $k \leq \ell$, we must fix an equipartition into ℓ columns on the x-axis and then find the optimal partition of the y-axis of size at most k . If we set the master partition Π of the OPTIMIZEXAXIS algorithm to be an equipartition into rows of size n , then it performs precisely the required optimization. Moreover, for fixed ℓ it can carry out the optimization simultaneously for all of the pairs $\{(2, \ell), \dots, (B/\ell, \ell)\}$ in time $O(|\Pi|^2 (B/\ell)\ell) = O(n^2 B)$. For fixed ℓ , this set contains all the pairs (k, ℓ) satisfying $k \leq \ell, k\ell \leq B$. Therefore, to compute all the required entries of $\widehat{M}(D)$ we need only apply this algorithm for each $\ell = 2, \dots, B/2$. Doing so gives a runtime of $O(n^2 B^2)$. \blacksquare

The algorithm above, while polynomial-time, is nonetheless not efficient enough for use in practice. However, a simple modification solves this problem without affecting the consistency of the resulting estimates. The modification hinges on the fact that OPTIMIZEXAXIS can use master partitions Π besides the equipartition of size n that we used above. Specifically, setting Π in the above algorithm to be an equipartition into ck “clumps”, where k is the size of the largest optimal partition being sought, speeds up the computation significantly. This modification gives a slightly different statistic, but one that has all of the theoretical properties of MIC_e —namely, consistent estimation of MIC_* and efficient exact computation. These properties are formalized in the following theorem.

Theorem 28 *Let (X, Y) be a pair of jointly distributed random variables, and let D_n be a sample of size n from the distribution of (X, Y) . For every $c \geq 1$, there exists a matrix $\{\widehat{M}\}^c(D_n)$ such that*

1. *The function*

$$\widetilde{\text{MIC}}_{e,B}(\cdot) = \max_{k\ell \leq B(n)} \{\widehat{M}\}^c(\cdot)_{k,\ell}$$

is a consistent estimator of MIC_ provided $\omega(1) < B(n) \leq O(n^{1-\varepsilon})$ for some $\varepsilon > 0$.*

2. *There exists an algorithm EQUICHARCLUMP for computing $r_B(\{\widehat{M}\}^c(D_n))$ in time $O(n + B^{5/2})$, which equals $O(n + n^{5(1-\varepsilon)/2})$ when $B(n) = O(n^{1-\varepsilon})$.*

Proof See Appendix H. ■

For an analysis of the effect of the parameter c in the above theorem on the results of the EQUICHARCLUMP algorithm, see Appendix H.3.

Setting $\varepsilon = 0.6$ in the above theorem yields the following corollary.

Corollary 29 *MIC_{*} can be estimated consistently in linear time.*

Of course, at low sample sizes, setting $\varepsilon = 0.6$ would be undesirable. However, our companion paper (Reshef et al., 2015a) shows empirically that at large sample sizes this strategy works very well on typical relationships.

We remark that the EQUICHARCLUMP algorithm given above is asymptotically faster even than the heuristic APPROX-MIC algorithm used to calculate MIC in practice, which runs in time $O(B(n)^4)$. As demonstrated in our companion paper (Reshef et al., 2015a), this difference translates into a substantial difference in runtimes for similar performance at a range of realistic sample sizes, ranging from a 30-fold speedup at $n = 500$ to over a 350-fold speedup at $n = 10,000$.

For readability, in the rest of this paper we do not distinguish between the two versions of MIC_e computed by the EQUICHAR and EQUICHARCLUMP algorithms described above. Wherever we present simulation data about MIC_e in simulations though, we use the version of the statistic computed by EQUICHARCLUMP.

4.4 Bias/Variance Characterization of MIC_e

The algorithm we presented in Section 3.5 for computing MIC_{*} to arbitrary precision in some cases allows us to examine the bias/variance properties of estimators of MIC_{*}. Here, we use it to examine the bias and variance of both MIC as computed by the heuristic APPROX-MIC algorithm from Reshef et al. (2011), and MIC_e as computed by the EQUICHARCLUMP algorithm given above. To do this, we performed a simulation analysis on the following set of relationships

$$\mathcal{Q} = \{(x + \varepsilon_\sigma, f(x) + \varepsilon'_\sigma) : x \in X_f, \varepsilon_\sigma, \varepsilon'_\sigma \sim \mathcal{N}(0, \sigma^2), f \in F, \sigma \in \mathbb{R}_{\geq 0}\}$$

where ε_σ and ε'_σ are i.i.d., F is the set of 16 functions analyzed in Reshef et al. (2011), and X_f is the set of n x-values that result in the points $(x_i, f(x_i))$ being equally spaced along the graph of f .

For each relationship $\mathcal{Z} \in \mathcal{Q}$ that we examined, we used the algorithm from Theorem 18 with very conservative values of k_0 and ℓ_0 to compute MIC_{*}. We then simulated 500 independent samples from \mathcal{Z} , each of size $n = 500$, and computed both APPROX-MIC and MIC_e on each one to obtain estimates of the sampling distributions of the two statistics. From each of the two sampling distributions, we estimated the bias and variance of either statistic on \mathcal{Z} . We then analyzed the bias, variance, and expected squared error of the two statistics as a function of relationship strength, which we quantified using the coefficient of determination (R^2) with respect to the generating function.

The results, presented in Figure 2, are interesting for two reasons. First, they demonstrate that for a typical usage parameter of $B(n) = n^{0.6}$, MIC_e performs substantially better than APPROX-MIC overall. Specifically, the median of the expected squared error of MIC_e

across the set F of functions is uniformly lower across R^2 values than that of APPROX-MIC. When average expected squared error is used instead of median, MIC_e still performs better on all but the strongest of relationships (R^2 above ~ 0.9). The superior performance of MIC_e is consistent with the fact that we have theoretical guarantees about its statistical properties whereas APPROX-MIC is a heuristic.

Second, the results show that different values of the exponent in $B(n) = n^\alpha$ give good performance in different signal-to-noise regimes due to a bias-variance trade-off represented by this parameter. We expand on this phenomenon and discuss its implications for choosing α in practice below.

4.4.1 CHOOSING $B(n)$

Large values of α lead to increased expected error in lower-signal regimes (low R^2) through both a positive bias in those regimes and a general increase in variance that predominantly affects those regimes. On the other hand, small values of α lead to an increased expected error in higher-signal regimes (high R^2) by leading to a negative bias in those regimes and by shifting the variance of the estimator toward those regimes. In other words, lower values of α are better suited for detecting weaker signals, while higher values of α are better suited for distinguishing among stronger signals. This is consistent with the results seen in our companion paper (Reshef et al., 2015a), which show that low values of α cause MIC_e to yield better powered independence tests while high values of α cause MIC_e to have better equitability.

Reshef et al. (2015a) provides simple, empirical recommendations about appropriate values of α for different settings. Those recommendations are formulated by choosing a set of representative relationships (e.g., a set of noisy functional relationships), as well as a “ground truth” population quantity Φ (e.g., R^2) that can be used to quantify the strength of each of those relationships, and then assessing which values of α maximize the equitability of MIC_e with respect to Φ at a given sample size. This approach is applied to an analysis of real data from the World Health Organization in Reshef et al. (2015a), and the parameters chosen for that analysis are the ones used for all subsequent analyses in this paper.

We remark that if the goal of the user is only detection of non-trivial relationships rather discovery of the strongest such relationships, α can also be chosen in a more straightforward manner: the user can subsample a small random set of relationships on which to compare the power of MIC_e for different values of α . Those relationships can then be discarded and the rest of the relationships analyzed with the optimal value of α . However, if the user’s primary goal is power against independence, the statistic TIC_e introduced in Section 5 of this paper should be used with this strategy rather than MIC_e .

4.5 Equitability of MIC_e

As mentioned previously, one of the main motivations for the introduction of MIC was equitability, the extent to which a measure of dependence usefully captures some notion of relationship strength on some set of standard relationships. We therefore carried out an empirical analysis of the equitability of MIC_e with respect to R^2 and compared its performance to distance correlation (Székely et al., 2007; Székely and Rizzo, 2009), mutual

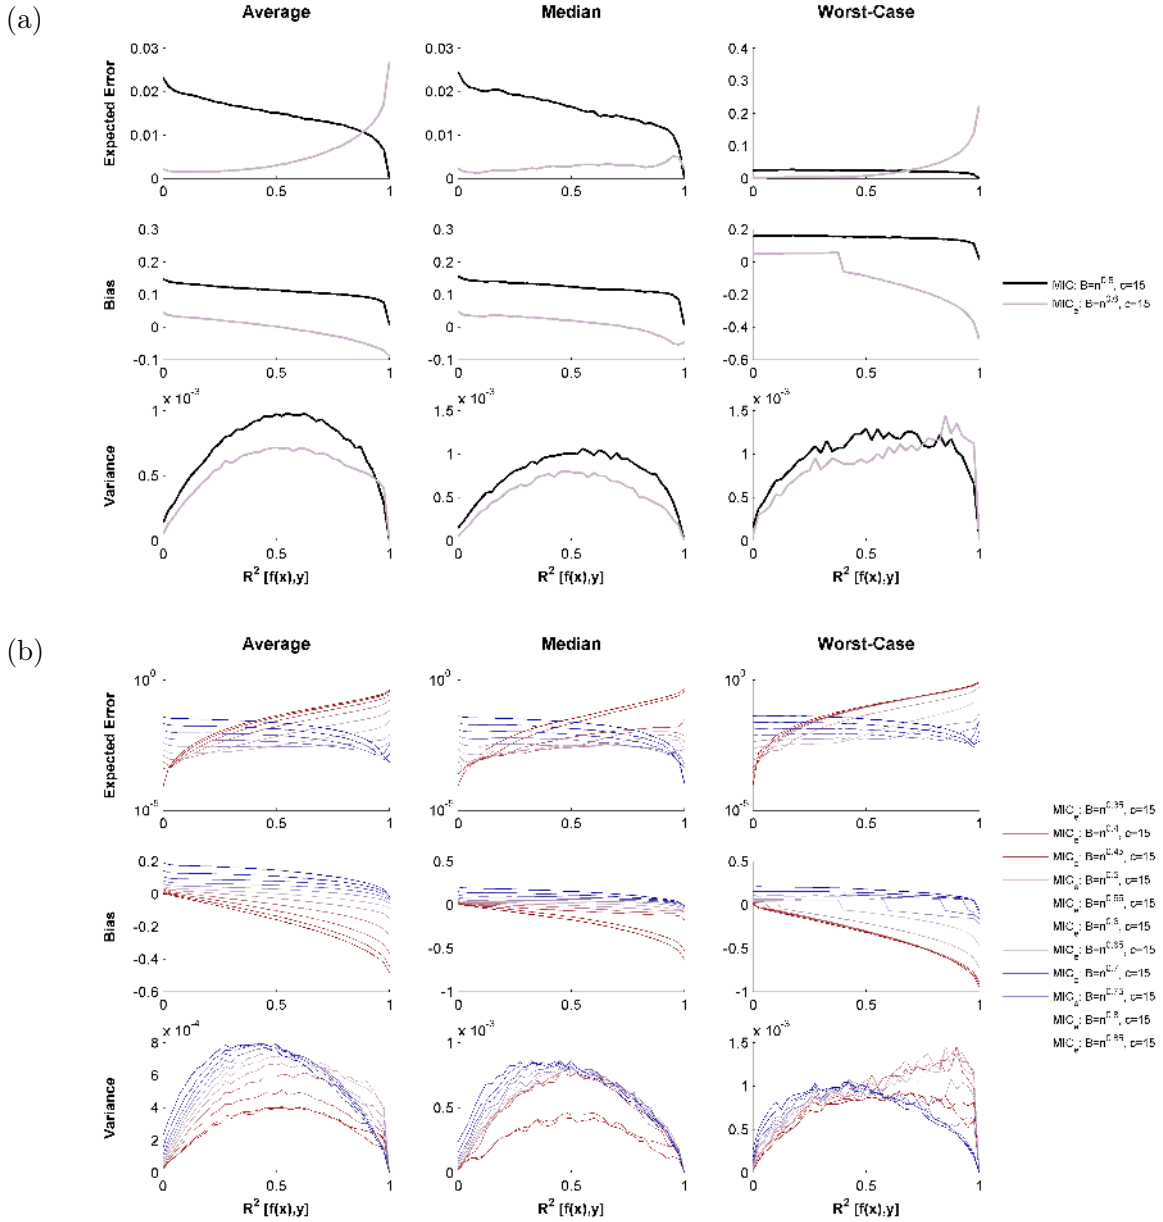


Figure 2: Bias/variance characterization of APPROX-MIC and MIC_e . Each plot shows expected squared error, bias, or variance across the set of noisy functional relationships described in Section 4.4 as a function of the R^2 of the relationships. The results are aggregated across the 16 function types analyzed by either the average, median, or worst result at every value of R^2 . (a) A comparison between MIC_e (light purple) and MIC as computed via the heuristic APPROX-MIC algorithm (black), at a typical usage parameter. (b) Performance of MIC_e with $B(n) = n^\alpha$ for various values of α .

information estimation (Kraskov et al., 2004), and maximal correlation estimation (Breiman and Friedman, 1985).

We began by assessing equitability on the set of relationships \mathcal{Q} defined above, a set that has been analyzed in previous work (Reshef et al., 2011, 2015a; Kinney and Atwal, 2014). The results, shown in Figure 3, confirm the superior equitability of the new estimator MIC_e on this set of relationships.

To assess equitability more objectively without relying on a manually curated set of functions, we then analyzed 160 random functions drawn from a Gaussian process distribution with a radial basis function kernel with one of eight possible bandwidths in the set $\{0.01, 0.025, 0.05, 0.1, 0.2, 0.25, 0.5, 1\}$ to represent a range of possible relationship complexities. The results, shown in Figure 4, show that MIC_e outperforms existing methods in terms of equitability with respect to R^2 on these functions as well. Appendix Figure J1 shows a version of this analysis under a different noise model that yields the same conclusion. We also examined the effect of outlier relationships on our results by repeatedly subsampling random subsets of 20 functions from this large set of relationships and measuring the equitability of each method on average over the subsets; results were similar.

One feature of the performance of MIC_e on these randomly chosen relationships that is demonstrated in Figure 4 is that it appears minimally sensitive to the bandwidth of the Gaussian process from which a given relationship is drawn. This puts it in contrast to, e.g., mutual information estimation, which shows a pronounced sensitivity to this parameter that prevents it from being highly equitable when relationships with different bandwidths are present in the same data set.

In our companion paper (Reshef et al., 2015a), we perform more in-depth analyses of the equitability with respect to R^2 of MIC_e , MIC, and the four measures of dependence described above as well as the Hilbert-Schmidt independence criterion (HSIC) (Gretton et al., 2005, 2007), the Heller-Heller-Gorfine (HHG) test (Heller et al., 2013), the data-derived partitions (DDP) test (Heller et al., 2016), and the randomized dependence coefficient (RDC) (Lopez-Paz et al., 2013). These analyses consider a range of sample sizes, noise models, marginal distributions, and parameter settings. They conclude that, in terms of equitability with respect to R^2 on the sets of noisy functional relationships studied, a) MIC_e uniformly outperforms MIC, and b) MIC_e outperforms all the methods tested in the large majority of settings examined. Appendix Figure I1 contains a reproduction of a representative equitability analysis from that paper for the reader’s reference.

5. The Total Information Coefficient

So far we have presented results about estimators of the population maximal information coefficient, a quantity for which equitability is the primary motivation. We now introduce and analyze a new measure of dependence, the *total information coefficient* (TIC). In contrast to the maximal information coefficient, the total information coefficient is designed not for equitability but rather as a test statistic for testing a null hypothesis of independence.

We begin by giving some intuition. Recall that the maximal information coefficient is the supremum of the characteristic matrix. While estimating the supremum of this matrix has many advantages, this estimation involves taking a maximum over many estimates of individual entries of the characteristic matrix. Since maxima of sets of random variables

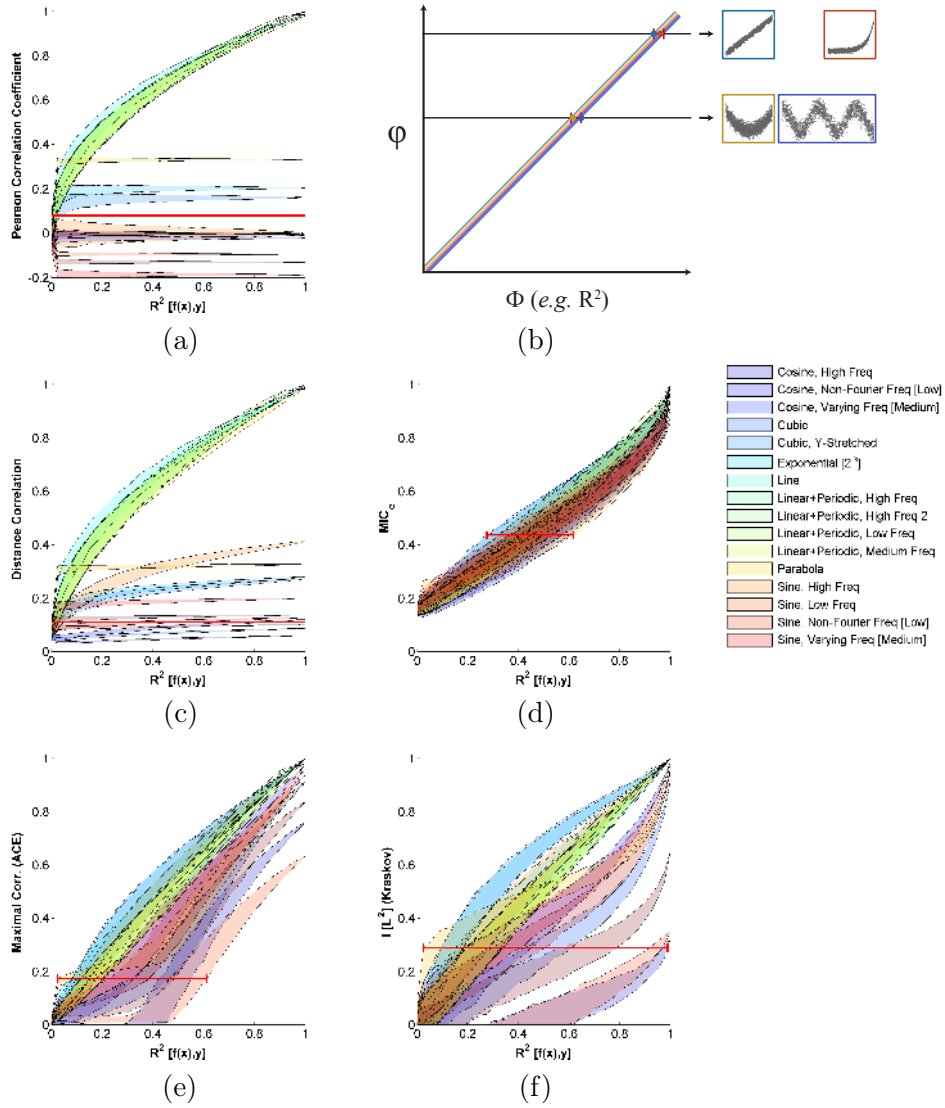


Figure 3: Equitability with respect to R^2 on a set of noisy functional relationships of (a) the Pearson correlation coefficient, (b) a hypothetical measure of dependence φ with perfect equitability, (c) distance correlation, (d) MIC_e , (e) maximal correlation estimation, and (f) mutual information estimation. For each relationship, a shaded region denotes estimated 5th and 95th percentile values of the sampling distribution of the statistic in question on that relationship at every R^2 . The resulting plot shows which values of R^2 correspond to a given value of each statistic. The red interval on each plot indicates the widest range of R^2 values corresponding to any one value of the statistic; the narrower the red interval, the higher the equitability. A red interval with width 0, as in (b), means that the statistic reflects only R^2 with no dependence on relationship type, as demonstrated by the pairs of thumbnails of relationships of different types with identical R^2 values.

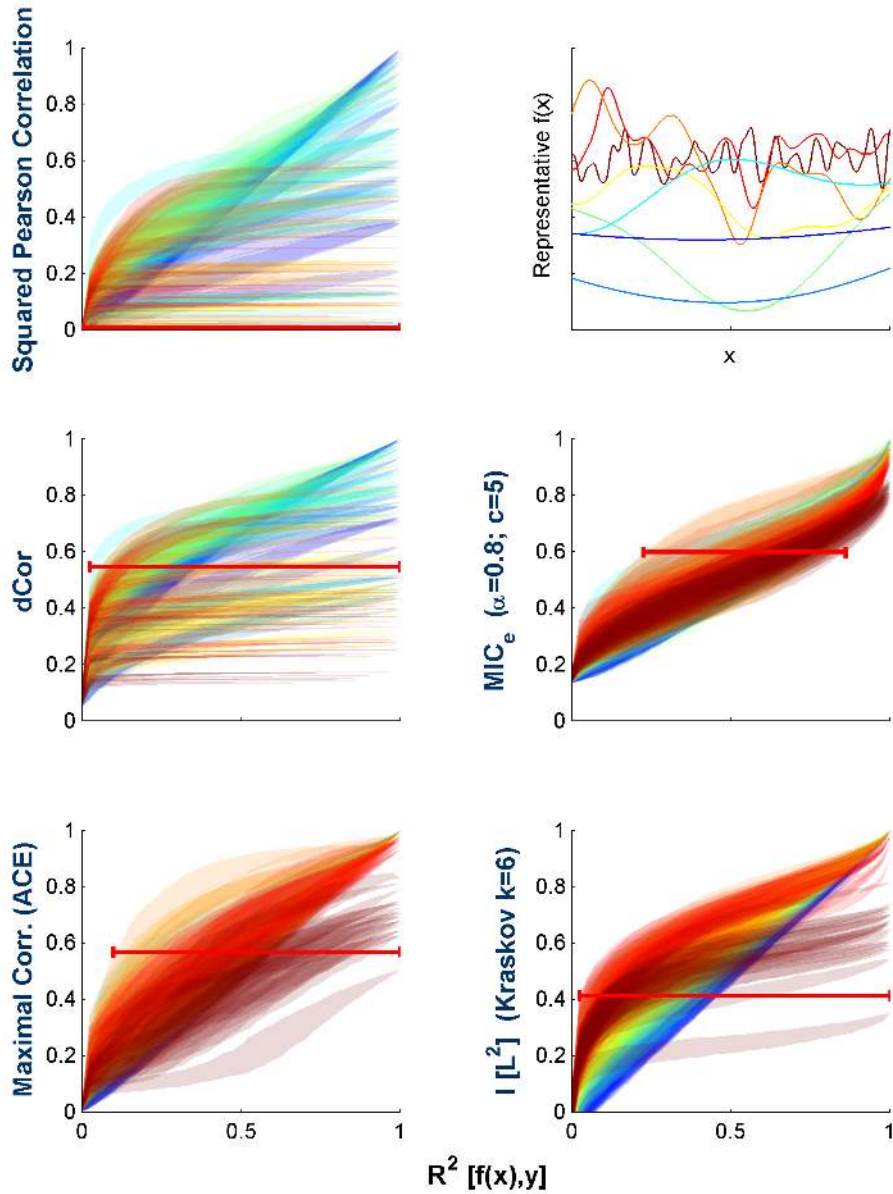


Figure 4: Equitability of methods examined on functions randomly drawn from a Gaussian process distribution. Each method is assessed as in Figure 3, with a red interval indicating the widest range of R^2 values corresponding to any one value of the statistic; the narrower the red interval, the higher the equitability. Each shaded region corresponds to one relationship, and the regions are colored by the bandwidth of the Gaussian process from which they were sampled. Sample relationships for each bandwidth are shown in the top right with matching colors.

tend to become large as the number of variables grows, one can imagine that this procedure may lead to an undesirable positive bias in the case of statistical independence, when the population characteristic matrix equals 0. This might be detrimental for independence testing, when the sampling distribution of a statistic under a null hypothesis of independence is crucial.

The intuition behind the total information coefficient is that if we instead consider a more stable property, such as the sum of the entries in the characteristic matrix, we might expect to obtain a statistic with a smaller bias in the case of independence and therefore better power. Stated differently, if our only goal is to distinguish any dependence at all from complete noise, then disregarding all of the sample characteristic matrix except for its maximal value may throw away useful signal, and the total information coefficient avoids this by summing all the entries.

We remark that in [Reshef et al. \(2011\)](#) it is suggested that other properties of the characteristic matrix may allow us to measure other aspects of a given relationship besides its strength, and several such properties were defined. The total information coefficient fits within this conceptual framework.

In this section we define the total information coefficient in the case of both the characteristic matrix (TIC) and the equicharacteristic matrix (TIC_e). We then prove that both TIC and TIC_e yield independence tests that are consistent against all dependent alternatives. (As in the case of MIC and MIC_e, TIC_e is more easily computable than TIC.) Finally, we present a simulation study of the power of independence testing based on TIC_e on an index set of relationships chosen in [Simon and Tibshirani \(2012\)](#), showing that TIC_e outperforms other common measures of dependence on many of the relationships and closely matches their performance on the rest.

5.1 Definition and Consistency of the Total Information Coefficient

We begin by defining the two versions of the total information coefficient. In the definition below, recall that \widehat{M} denotes a sample characteristic matrix whereas $[\widehat{M}]$ denotes a sample equicharacteristic matrix.

Definition 30 *Let $D \subset \mathbb{R}^2$ be a set of n ordered pairs, and let $B : \mathbb{Z}^+ \rightarrow \mathbb{Z}^+$. We define*

$$TIC_B(D) = \sum_{k\ell \leq B(n)} \widehat{M}(D)_{k,\ell}$$

and

$$TIC_{e,B}(D) = \sum_{k\ell \leq B(n)} [\widehat{M}](D)_{k,\ell}.$$

To show that these two statistics lead to consistent independence tests, we must take a step back and analyze the behavior of the analogous population quantities.

Definition 31 *For a matrix A and a positive number B , the B -partial sum of A , denoted by $S_B(A)$, is*

$$S_B(A) = \sum_{k\ell \leq B} A_{k,\ell}.$$

When A is an (equi)characteristic matrix, $S_B(A)$ is the sum over all entries corresponding to grids with at most B total cells. Thus, if $\widehat{M}(D)$ is a sample characteristic matrix of a sample D , $S_B(\widehat{M}(D)) = \text{TIC}_B(D)$, and the same holds for $S_B([\widehat{M}](D))$ and $\text{TIC}_{e,B}(D)$.

It is clear that if X and Y are statistically independent random variables, then both the characteristic matrix $M(X, Y)$ and the equicharacteristic matrix $[M](X, Y)$ are identically 0, so that $S_B(M(X, Y)) = S_B([M](X, Y)) = 0$ for all B . However, we are also interested in how these quantities behave when X and Y are dependent. The following pair of propositions helps us understand this. The first proposition shows a lower bound on the values of entries in both $M(X, Y)$ and $[M](X, Y)$. The second proposition translates this into an asymptotic characterization of how quickly $S_B(M)$ and $S_B([M])$ grow as functions of B . These two propositions are the technical heart of why the total information coefficient yields a consistent independence test.

Proposition 32 *Let (X, Y) be a pair of jointly distributed random variables. If X and Y are statistically independent, then $M(X, Y) \equiv [M](X, Y) \equiv 0$. If not, then there exists some $a > 0$ and some integer $\ell_0 \geq 2$ such that*

$$M(X, Y)_{k,\ell}, [M](X, Y)_{k,\ell} \geq \frac{a}{\log \min\{k, \ell\}}$$

either for all $k \geq \ell \geq \ell_0$, or for all $\ell \geq k \geq \ell_0$.

Proof See Appendix K.1 ■

Proposition 33 *Let (X, Y) be a pair of jointly distributed random variables. If X and Y are statistically independent, then $S_B(M(X, Y)) = S_B([M](X, Y)) = 0$ for all $B > 0$. If not, then $S_B(M(X, Y))$ and $S_B([M](X, Y))$ are both $\Omega(B \log \log B)$.*

Proof See Appendix K.2 ■

The propositions above, together with reasoning analogous to the convergence arguments presented earlier, can be used to show the main result of this section, namely that the statistics TIC and TIC_e yield consistent independence tests.

Theorem 34 *The statistics TIC_B and $\text{TIC}_{e,B}$ yield consistent right-tailed tests of independence, provided $\omega(1) < B(n) \leq O(n^{1-\varepsilon})$ for some $\varepsilon > 0$.*

Proof See Appendix K.3. ■

In practice, we often use the EQUICHARCLUMP algorithm (see Section 4.3) to compute the equicharacteristic matrix from which we calculate TIC_e . This algorithm does not compute the sample equicharacteristic matrix exactly. However, as in the case of MIC_e , the use of the algorithm does not affect the theoretical properties of the statistic. This is proven in Appendix H.

5.2 Power of Independence Tests Based on TIC_e

With the consistency of independence tests based on TIC and TIC_e established, we turn now to empirical evaluation of the power of independence testing based on TIC_e as computed using the EQUICHARCLUMP algorithm.

To evaluate the power of TIC_e -based tests, we reproduced the analysis performed in [Simon and Tibshirani \(2012\)](#). Namely, we considered the set of relationships they analyzed, defined by

$$\mathcal{Q} = \{(X, f(X) + \varepsilon') : X \sim \text{Unif}, f \in F, \varepsilon' \sim \mathcal{N}(0, \sigma^2), \sigma \in \mathbb{R}_{\geq 0}\}.$$

where F is a set of functions specified in [Simon and Tibshirani \(2012\)](#). (NB: one of the relationships is a circle, which we treat as a union of two half-circles.)

For each relationship \mathcal{Z} in this set that we examined, we simulated a null hypothesis of independence with the same marginal distributions, and generated 1,000 independent samples, each with a sample size of $n = 500$, from both \mathcal{Z} and from the null distribution. These were used to estimate the power of the size- α right-tailed independence test based on each statistic being evaluated. Following Simon and Tibshirani, we compared TIC_e to the distance correlation ([Székely et al., 2007](#); [Székely and Rizzo, 2009](#)), the original maximal information coefficient ([Reshef et al., 2011](#)) as approximated using APPROX-MIC, and to the Pearson correlation. (Though it is not a measure of dependence, the Pearson correlation was presumably included by Simon and Tibshirani as an intuitive benchmark for what is achievable under a linear model.) We also compared to MIC_e using identical parameters to those of TIC_e to examine whether the summation performed by TIC_e is better than maximization when all other things are equal. Note that we do not compare to methods of analyzing contingency tables, such as Pearson’s chi-squared test. This is because our data are real-valued rather than discrete, and so contingency-based methods are not applicable. However, when data are discrete, those methods can be very well powered.

The results of our analysis are presented in [Figure 5](#). First, the figure shows that TIC_e compares quite favorably with distance correlation, a method considered to have state-of-the-art power ([Simon and Tibshirani, 2012](#)). Specifically, TIC_e uniformly outperforms distance correlation on 5 of the 8 relationship types examined, and performs comparably to it on the other three relationship types. We remark that distance correlation has many advantages over TIC_e , including the fact that it easily generalizes to higher-dimensional relationships and comes with an elegant and comprehensive theoretical framework.

The analysis also shows that TIC_e outperforms the original maximal information coefficient by a very large margin, and outperforms MIC_e as well, supporting the intuition that the summation performed by the former can indeed lead to substantial gains in power against independence over the maximization performed by the latter. (We note that in both Simon and Tibshirani’s analysis and in this one, the original maximal information coefficient was run with default parameters that were optimized for equitability rather than power against independence. When run with different parameters, its power improves substantially, though it still does not match the power of MIC_e . See [Appendix Figure I2](#) and the discussion in [Reshef et al., 2015a](#).)

Our companion paper ([Reshef et al., 2015a](#)) expands on this analysis, conducting an in-depth evaluation of the the power against independence of the tests described above as well

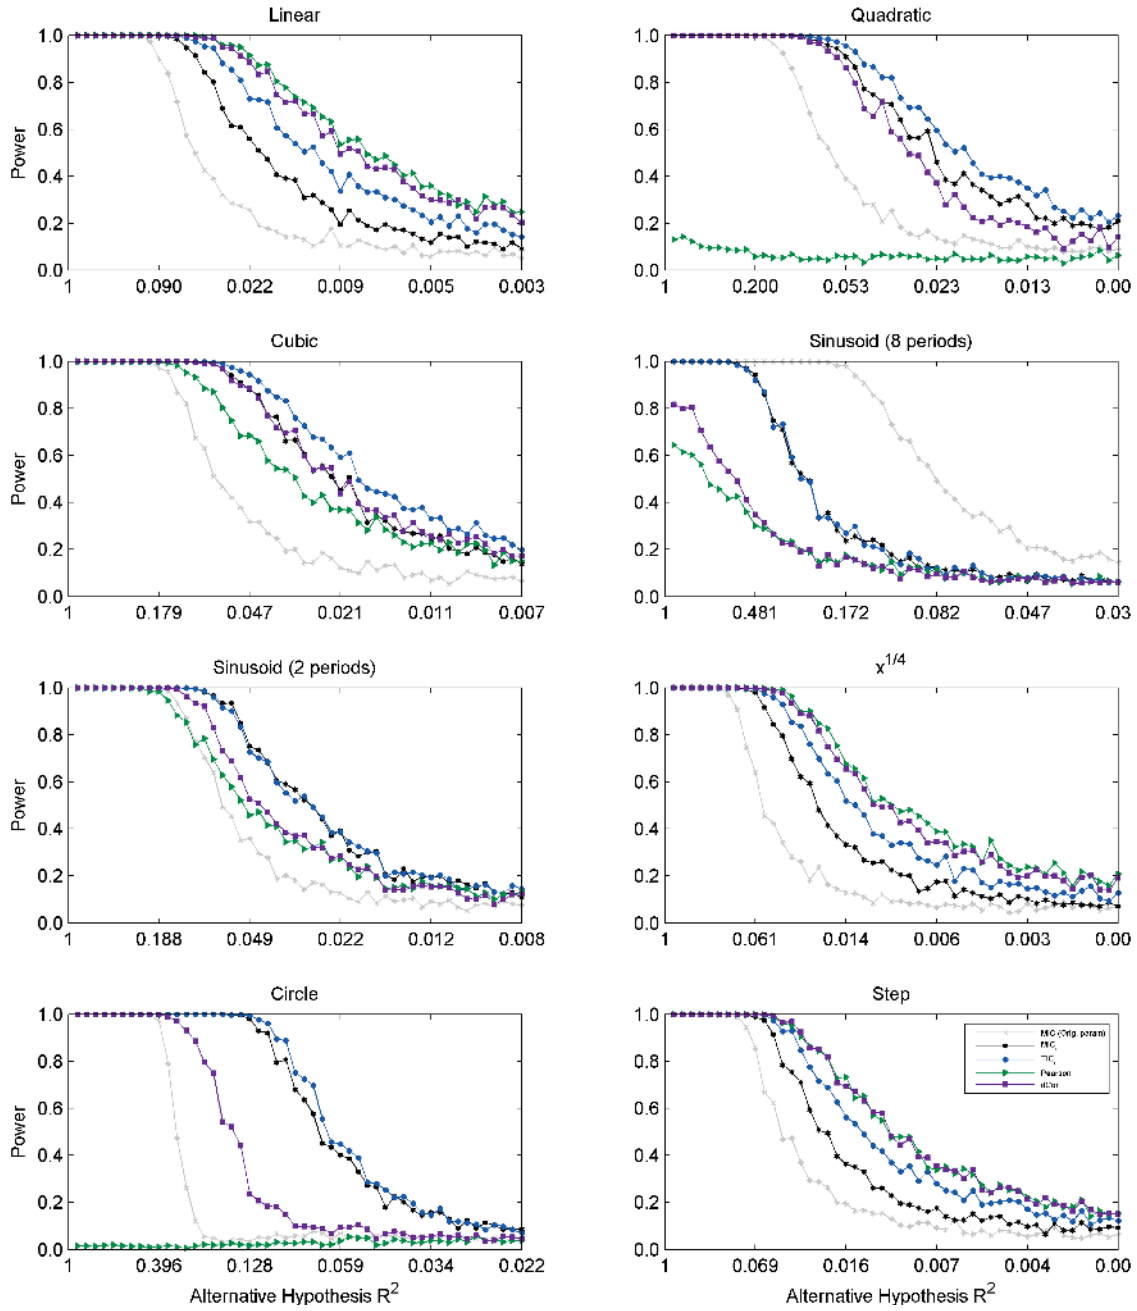


Figure 5: Comparison of power of independence testing based on TIC_e (blue) to MIC with default parameters (gray), MIC_e with the same parameters as TIC_e (black), distance correlation (purple), and the Pearson correlation coefficient (green) across several alternative hypothesis relationship types chosen by Simon and Tibshirani (2012). The relationships analyzed are described in Section 5.2.

as tests based on mutual information estimation (Kraskov et al., 2004), maximal correlation estimation (Breiman and Friedman, 1985), HSIC (Gretton et al., 2005, 2007), HHG (Heller et al., 2013), DDP (Heller et al., 2016), and RDC (Lopez-Paz et al., 2013). These analyses consider a range of sample sizes and parameter settings, as well as a variety of ways of quantifying power across different alternative hypothesis relationship types and noise levels. They conclude that in most settings TIC_e either outperforms all the methods tested or performs comparably to the best ones. Appendix Figure I2 contains a reproduction of one detailed set of power curves from the main analysis in that paper for the reader’s reference.

6. Conclusion

As high-dimensional data sets become increasingly common, data exploration requires not only statistics that can accurately detect a large number of non-trivial relationships in a data set, but also ones that can identify a smaller number of strongest relationships. The former property is achieved by measures of dependence that yield independence tests with high power; the latter is achieved by measures of dependence that are equitable with respect to some measure of relationship strength. In this paper, we introduced two related measures of dependence that achieve these two goals, respectively, through the following theoretical contributions.

- A new population measure of dependence, MIC_* , that we proved can be viewed in three different ways: as the population value of the maximal information coefficient (MIC) from Reshef et al. (2011), as a “minimal smoothing” of mutual information that makes it uniformly continuous, or as the supremum of an infinite sequence defined in terms of optimal partitions of one marginal at a time of a given joint distribution.
- An efficient approach for approximating the MIC_* of a given joint distribution.
- A statistic MIC_e that is a consistent estimator of MIC_* , is efficiently computable, and has good equitability with respect to R^2 both on a manually chosen set of noisy functional relationships as well as on a set of randomly chosen noisy functional relationships.
- The total information coefficient (TIC_e), a statistic that arises as a trivial side-product of the computation of MIC_e and yields a consistent and powerful independence test.

Though we presented here some empirical results for MIC_* , MIC_e , and TIC_e , our focus was on theoretical considerations; the performance of these methods is analyzed in detail in our companion paper (Reshef et al., 2015a). That paper shows that on a large set of noisy functional relationships with varying noise and sampling properties, the asymptotic equitability with respect to R^2 of MIC_* is quite high and the equitability with respect to R^2 of MIC_e is state-of-the-art. It also shows that the power of the independence test based on TIC_e is state-of-the-art across a wide variety of dependent alternative hypotheses. Finally, it demonstrates that the algorithms presented here allow for MIC_e and TIC_e to be computed simultaneously very quickly, enabling analysis of extremely large data sets using both statistics together.

Our contributions are of both theoretical and practical importance for several reasons. First, our characterization of MIC_* as the large-sample limit of MIC sheds light on the latter statistic. For example, while MIC is parametrized, MIC_* is not. Knowing that MIC converges in probability to MIC_* tells us that this parametrization is statistical only: it controls the bias/variance properties of the statistic, but not its asymptotic behavior.

Second, the normalization in the definition of MIC , while empirically seen to yield good performance, had previously not been theoretically understood. Our result that this normalization is the minimal smoothing necessary to make mutual information uniformly continuous provides for the first time a lens through which the normalization is canonical. In doing so, it constitutes an initial step toward understanding the role of the normalization in the performance of MIC_* and MIC . The uniform continuity of MIC_* and the lack of continuity of ordinary mutual information also suggest that estimation of the former may be easier in some sense than estimation of the latter. This is consonant with a recent result concerning difficulty of estimation of mutual information shown in [Ding and Li \(2013\)](#). It is also borne out empirically by the substantial finite-sample bias and variance observed in [Reshef et al. \(2015a\)](#) of the Kraskov mutual information estimator ([Kraskov et al., 2004](#)) compared to MIC_e .

Third, our alternate characterization of MIC_* in terms of one-dimensional optimization over partitions rather than two-dimensional optimization over grids enhances our understanding of how to efficiently compute it in the large-sample limit and estimate it from finite samples using MIC_e . This is a significant improvement over the previous state of affairs, in which the statistic MIC could only be approximated heuristically, with even the heuristic approximation being orders of magnitude slower than the results in this paper now allow.

Finally, the introduction of the total information coefficient provides evidence that the basic approach of considering the set of normalized mutual information values achievable by applying different grids to a joint distribution is of fundamental value in characterizing dependence. Interestingly, a statistic introduced in [Heller et al. \(2016\)](#) follows a similar approach by considering the (non-normalized) sum of the mutual information values achieved by all possible finite grids. Consistent with our demonstration here that an aggregative grid-based approach works well, that statistic also achieves excellent power. (TIC_e is compared to the statistic from [Heller et al. 2016](#) in our companion paper, [Reshef et al., 2015a](#).)

Taken together, our results point to joint use of the statistics MIC_e and TIC_e as a theoretically grounded, computationally efficient, and highly practical approach to data exploration. Specifically, since the two statistics can be computed simultaneously with little extra cost beyond that of computing either individually, we propose computing both of them on all variable pairs in a data set, using TIC_e to filter out non-significant associations, and then using MIC_e to rank the remaining variable pairs. Such a strategy would have the advantage of leveraging the state-of-the-art power of TIC_e to substantially reduce the multiple-testing burden on MIC_e , while utilizing the latter statistic’s state-of-the-art equitability to effectively rank relationships for follow-up by the practitioner.

Our results, while useful, nevertheless have limitations that warrant exploration in future work. First, for a sample D from the distribution of some random (X, Y) , all of the sample quantities we define here use the naive estimate $I(D|_G)$ of the quantity $I((X, Y)|_G)$ for various grids G . There is a long and fruitful line of work on more sophisticated estimators of the discrete mutual information [Paninski \(2003\)](#) whose use instead of $I(D|_G)$ could improve

the statistics introduced here. Second, our approach to approximating the MIC_* of a given joint density consists of computing a finite subset of an infinite set whose supremum we seek to calculate. However, the choice of how large a finite set we should compute in order to approximate the supremum to a given precision remains heuristic. Finally, though empirical characterization of the equitability of MIC_e on representative sets of relationships is important and promising, we are still missing a theoretical characterization of its equitability in the large-sample limit. A clear theoretical demarcation of the set of relationships on which MIC_* achieves good equitability with respect to R^2 , and an understanding of why that is, would greatly advance our understanding of both MIC_* and equitability.

Acknowledgments

We would like to acknowledge R Adams, E Airoidi, T Broderick, A Gelman, M Gorfine, R Heller, J Huggins, T Jaakkola, J Mueller, J Tenenbaum, and R Tibshirani for constructive conversations and useful feedback. HKF was supported by the Fannie and John Hertz Foundation. MM was supported in part by NSF grants CCF-1563710, CCF-1535795, CCF-1320231, and CNS-1228598. DNR and YAR were supported by the Paul and Daisy Soros Foundation. YAR was supported by award Number T32GM007753 from the National Institute of General Medical Sciences, as well as the National Defense Science and Engineering Graduate Fellowship. PCS was supported by the Howard Hughes Medical Institute. The content of this paper is solely the responsibility of the authors and does not necessarily represent the official views of the National Institute of General Medical Sciences or the National Institutes of Health.

Appendix A. Proof of Theorem 6

This appendix is devoted to proving Theorem 6, restated below.

Theorem *Let $f : m^\infty \rightarrow \mathbb{R}$ be uniformly continuous, and assume that $f \circ r_i \rightarrow f$ pointwise. Then for every random variable (X, Y) , we have*

$$(f \circ r_{B(n)}) \left(\widehat{M}(D_n) \right) \rightarrow f(M(X, Y))$$

in probability where D_n is a sample of size n from the distribution of (X, Y) , provided $\omega(1) < B(n) \leq O(n^{1-\varepsilon})$ for some $\varepsilon > 0$.

We prove the theorem by a sequence of lemmas that build on each other to bound the bias of $I^*(D, k, \ell)$. The general strategy is to capture the dependencies between different k -by- ℓ grids G by considering a “master grid” Γ that contains many more than $k\ell$ cells. Given this master grid, we first bound the difference between $I(D|_G)$ and $I((X, Y)|_G)$ only for sub-grids G of Γ . The bound is in terms of the difference between $D|_\Gamma$ and $(X, Y)|_\Gamma$. We then show that this bound can be extended without too much loss to all k -by- ℓ grids. This gives what we seek, because then the difference between $I(D|_G)$ and $I((X, Y)|_G)$ is uniformly bounded for all grids G in terms of the same random variable: $D|_\Gamma$. Once this is done, standard arguments give the consistency we seek.

In our argument we occasionally require technical facts about entropy and mutual information that are self-contained and unrelated to the central ideas. These lemmas are consolidated in Appendix L.

We begin by using one of these technical lemmas to prove a bound on the difference between $I(D|_G)$ and $I((X, Y)|_G)$ that is uniform over all grids G that are sub-grids of a much denser grid Γ . The common structure imposed by Γ will allow us to capture the dependence between the quantities $|I(D|_G) - I((X, Y)|_G)|$ for different grids G .

Lemma 35 *Let $\Pi = (\Pi_X, \Pi_Y)$ and $\Psi = (\Psi_X, \Psi_Y)$ be random variables distributed over the cells of a grid Γ , and let $(\pi_{i,j})$ and $(\psi_{i,j})$ be their respective distributions. Define*

$$\varepsilon_{i,j} = \frac{\psi_{i,j} - \pi_{i,j}}{\pi_{i,j}}.$$

Let G be a sub-grid of Γ with B cells. Then for every fixed $0 < a < 1$ we have

$$|I(\Psi|_G) - I(\Pi|_G)| \leq O\left((\log B) \sum_{i,j} |\varepsilon_{i,j}|\right)$$

when $|\varepsilon_{i,j}| \leq 1 - a$ for all i and j .

Proof Let $P = \Pi|_G$ and $Q = \Psi|_G$ be the random variables induced by Π and Ψ respectively on the cells of G . Using the fact that $I(X, Y) = H(X) + H(Y) - H(X, Y)$, we write

$$|I(Q) - I(P)| \leq |H(Q_X) - H(P_X)| + |H(Q_Y) - H(P_Y)| + |H(Q) - H(P)|$$

where Q_X and P_X denote the marginal distributions on the columns of G and Q_Y and P_Y denote the marginal distributions on the rows. We can bound each of the terms on the right-hand side of the equation above using a Taylor expansion argument given in Lemma 51, whose proof is found in Appendix L. Doing so gives

$$|I(Q) - I(P)| \leq (\ln B) \left(\sum_i O(|\varepsilon_{i,*}|) + \sum_j O(|\varepsilon_{*,j}|) + \sum_{i,j} O(|\varepsilon_{i,j}|) \right)$$

where

$$\varepsilon_{i,*} = \frac{\sum_j (\psi_{i,j} - \pi_{i,j})}{\sum_j \pi_{i,j}}$$

and $\varepsilon_{*,j}$ is defined analogously.

To obtain the result, we observe that

$$|\varepsilon_{i,*}| = \left| \frac{\sum_j \pi_{i,j} \varepsilon_{i,j}}{\sum_j \pi_{i,j}} \right| \leq \frac{\sum_j \pi_{i,j} |\varepsilon_{i,j}|}{\sum_j \pi_{i,j}} \leq \sum_j |\varepsilon_{i,j}|$$

since $\pi_{i,j} / \sum_j \pi_{i,j} \leq 1$, and the analogous bound holds for $|\varepsilon_{*,j}|$. ■

We now extend Lemma 35 to all grids with B cells rather than just those that are sub-grids of the master grid Γ . The proof of this lemma relies on an information-theoretic result proven in Appendix B that bounds the difference in mutual information between two distributions that can be obtained from each other by moving a small amount of probability mass.

Lemma 36 *Let $\Pi = (\Pi_X, \Pi_Y)$ and $\Psi = (\Psi_X, \Psi_Y)$ be random variables, and let Γ be a grid. Define $\varepsilon_{i,j}$ on $\Pi|_\Gamma$ and $\Psi|_\Gamma$ as in Lemma 35. Let G be any grid with B cells, and let δ (resp. d) represent the total probability mass of $\Pi|_\Gamma$ (resp. $\Psi|_\Gamma$) falling in cells of Γ that are not contained in individual cells of G . We have that*

$$|I(\Psi|_G) - I(\Pi|_G)| \leq O\left(\left(\sum_{i,j} |\varepsilon_{i,j}| + \delta + d\right) \log B + \delta \log(1/\delta) + d \log(1/d)\right)$$

provided that the $|\varepsilon_{i,j}|$ are bounded away from 1 and that $d, \delta \leq 1/2$.

Proof In the proof below, we use the convention that for any two grids G and G' and any random variable Z , the expression $\Delta^Z(G, G')$ denotes $|I(Z|_G) - I(Z|_{G'})|$.

Consider the grid G' obtained by replacing every horizontal or vertical line in G that is not in Γ with a closest line in Γ . The grid G' is clearly a sub-grid of Γ . Moreover, $\Pi|_{G'}$ (resp. $\Psi|_{G'}$) can be obtained from $\Pi|_G$ (resp. $\Pi|_G$) by moving at most δ (resp. d) probability mass. This can be shown to imply that

$$\Delta^\Pi(G, G') \leq O(\delta \log(1/\delta) + \delta \log B) \quad \text{and} \quad \Delta^\Psi(G', G) \leq O(d \log(1/d) + d \log B).$$

The proof of this information-theoretic fact is self-contained and so we defer it to Proposition 40 in Appendix B, as it is more central to the arguments presented there.

With $\Delta^\Phi(G, G')$ and $\Delta^\Psi(G', G)$ bounded in terms of δ and d , we can bound $|I(\Psi|_G) - I(\Phi|_G)|$ using the triangle inequality by comparing it with

$$\Delta^\Pi(G, G') + |I(\Pi|_{G'}) - I(\Psi|_{G'})| + \Delta^\Psi(G', G)$$

and bounding the middle term using Lemma 35, since $G' \subset \Gamma$. ■

We now use the fact that the variables $\varepsilon_{i,j}$ defined in Lemma 35 are small with high probability to give a concrete bound on the bias of $I(D|_G)$ that is uniform over all k -by- ℓ grids G and that holds with high probability. It is useful at this point to recall that, given a distribution (X, Y) , an *equipartition* of (X, Y) is a grid G such that all the rows of $(X, Y)|_G$ have the same probability mass, and all the columns do as well.

Lemma 37 *Let D_n be a sample of size n from the distribution of a pair (X, Y) of jointly distributed random variables. For any $\alpha \geq 0$, any $\varepsilon > 0$, and any integers $k, \ell > 1$, we have that for all n*

$$|I(D_n|_G) - I((X, Y)|_G)| \leq O\left(\frac{\log(k\ell)}{C(n)^\alpha} + \frac{\log(k\ell n)}{n^{\varepsilon/4}}\right)$$

for every k -by- ℓ grid G with probability at least $1 - C(n)e^{-\Omega(n/C(n)^{1+2\alpha})}$, where $C(n) = k\ell n^{\varepsilon/2}$.

Proof Fix n , and let Γ be an equipartition of (X, Y) into $kn^{\varepsilon/4}$ rows and $\ell n^{\varepsilon/4}$ columns. $C(n)$ is now the number of cells in Γ . Lemma 36, with $\Pi = (X, Y)$ and $\Psi = D$, shows that $|I(D|_{\Gamma}) - I((X, Y)|_{\Gamma})|$ is at most

$$O\left(\left(\sum_{i,j} |\varepsilon_{i,j}| + \delta + d\right) \log(k\ell) + \delta \log(1/\delta) + d \log(1/d)\right)$$

provided the $\varepsilon_{i,j}$ have absolute value bounded away from 1, and provided that $d, \delta \leq 1/2$.

The remainder of the proof proceeds as follows. We first show that the $\varepsilon_{i,j}$ are small with high probability. This will both show that the lemma's requirement on the $\varepsilon_{i,j}$ holds and allow us to bound the sum in the inequality above. We will then use our bound on the $\varepsilon_{i,j}$ to bound d in terms of δ . Finally, we will bound δ using the fact that the number of rows and columns in Γ increases with n . This will give us that $d, \delta \leq 1/2$ and allow us to bound the rest of the terms in the expression above.

Bounding the $\varepsilon_{i,j}$: We bound the $\varepsilon_{i,j}$ using a multiplicative Chernoff bound. Let $\pi_{i,j}$ and $\psi_{i,j}$ represent the probability mass functions of $(X, Y)|_{\Gamma}$ and $D|_{\Gamma}$ respectively. We write

$$\begin{aligned} \mathbf{P}(|\varepsilon_{i,j}| \geq \delta) &= \mathbf{P}(\pi_{i,j}(1 - \delta) \leq \psi_{i,j} \leq \pi_{i,j}(1 + \delta)) \\ &\leq e^{-\Omega(n\pi_{i,j}\delta^2)} \end{aligned}$$

since $\psi_{i,j}$ is a sum of n i.i.d. Bernoulli random variables and $\mathbf{E}(\psi_{i,j}) = n\pi_{i,j}$. (See, e.g., [Mitzenmacher and Upfal 2005](#).) Setting $\delta = \sqrt{\pi_{i,j}}/C(n)^{1/2+\alpha}$ yields

$$\mathbf{P}\left(|\varepsilon_{i,j}| \geq \frac{\sqrt{\pi_{i,j}}}{C(n)^{1/2+\alpha}}\right) \leq e^{-\Omega(n/C(n)^{1+2\alpha})}.$$

A union bound over the pairs (i, j) then gives that, with the desired probability, the above bound on $|\varepsilon_{i,j}|$ holds for all i, j .

Bounding $\sum |\varepsilon_{i,j}|$: The bound on the $\varepsilon_{i,j}$ implies that

$$\begin{aligned} \sum_i |\varepsilon_{i,j}| &\leq \frac{1}{C(n)^{1/2+\alpha}} \sum_{i,j} \sqrt{\pi_{i,j}} \\ &\leq \frac{1}{C(n)^{1/2+\alpha}} \sqrt{C(n)} \\ &\leq \frac{1}{C(n)^\alpha} \end{aligned}$$

where the second line follows from the fact that the function $\sum \sqrt{\pi_{i,j}}$ is symmetric and concave and therefore, when restricted to the hyperplane $\sum \pi_{i,j} = 1$, must achieve its maximum when $\pi_{i,j} = 1/C(n)$ for all i, j .

Bounding d in terms of δ : We use our bound on the $\varepsilon_{i,j}$ to bound d . We do so by observing that it implies

$$\psi_{i,j} \leq \pi_{i,j} \left(1 + \frac{\sqrt{\pi_{i,j}}}{C(n)^{1/2+\alpha}}\right) = \pi_{i,j} + \frac{\pi_{i,j}^{3/2}}{C(n)^{1/2+\alpha}} \leq \pi_{i,j} + \frac{\pi_{i,j}}{C(n)^{1/2+\alpha}} \leq 2\pi_{i,j}$$

since $\pi_{i,j} \leq 1$ and $C(n) \geq 1$.

The connection to d comes from the fact that for any column j of Γ , this means that

$$\psi_{*,j} = \sum_i \psi_{i,j} \leq 2 \sum_i \pi_{i,j} = 2\pi_{*,j}.$$

This also applies to the sums across rows. Since d is a sum of terms of the form $\psi_{*,j}$ and $\psi_{i,*}$ for j in some index set J and i in an index set I , and δ is a sum of terms of the form $\pi_{*,j}$ and $\pi_{i,*}$ with the same index sets, we therefore get that $d \leq 2\delta$.

Bounding δ and obtaining the result: To bound δ , we observe that because G has at most $\ell - 1$ vertical lines and $k - 1$ horizontal lines, we have

$$\delta \leq \frac{\ell}{\ell n^{\varepsilon/4}} + \frac{k}{k n^{\varepsilon/4}} \leq \frac{2}{n^{\varepsilon/4}}.$$

This bound on δ allows us to bound the terms involving d and δ by

$$\delta + d \leq O\left(\frac{1}{n^{\varepsilon/4}}\right), \quad \delta \log\left(\frac{1}{\delta}\right) + d \log\left(\frac{1}{d}\right) \leq O\left(\frac{\log n}{n^{\varepsilon/4}}\right).$$

Combining all of the bounds gives the desired result. ■

Our final lemma shows that as long as $B(n)$ doesn't grow too fast, the bound from the previous lemma yields a uniform bound on the entire sample characteristic matrix. This is done by specifying an error threshold for which Lemma 37 yields a bound that holds with high probability, and then invoking a union bound.

Lemma 38 *Let D_n be a sample of size n from the distribution of a pair (X, Y) of jointly distributed random variables. For every $B(n) = O(n^{1-\varepsilon})$, there exists an $a > 0$ such that for sufficiently large n ,*

$$\left| \widehat{M}(D_n)_{k,\ell} - M(X, Y)_{k,\ell} \right| \leq O\left(\frac{1}{n^a}\right)$$

holds for all $k\ell \leq B(n)$ with probability $P(n) = 1 - o(1)$, where $\widehat{M}(D_n)_{k,\ell}$ is the k, ℓ -th entry of the sample characteristic matrix and $M(X, Y)_{k,\ell}$ is the k, ℓ -th entry of the population characteristic matrix of (X, Y) .

Proof Fix k, ℓ , and any α satisfying $0 < \alpha < \varepsilon/(4 - 2\varepsilon)$. Lemma 37 implies that with high probability the difference $|\widehat{M}(D_n)_{k,\ell} - M_{k,\ell}|$ is at most

$$\begin{aligned} O\left(\frac{\log(k\ell)}{C(n)^\alpha} + \frac{\log(k\ell n)}{n^{\varepsilon/4}}\right) &\leq O\left(\frac{\log n}{C(n)^\alpha} + \frac{\log n}{n^{\varepsilon/4}}\right) \\ &\leq O\left(\frac{\log n}{n^{\alpha\varepsilon/2}} + \frac{\log n}{n^{\varepsilon/4}}\right) \end{aligned}$$

where the first inequality comes from $k\ell \leq B(n)$ and second is because $C(n) = k\ell n^{\varepsilon/2} \geq n^{\varepsilon/2}$. This bound is at most $O(1/n^a)$ for every $a < \min\{\alpha\varepsilon/2, \varepsilon/4\}$, as desired. It remains only to show that the bound holds with high probability across all $k\ell \leq B(n)$.

Lemma 37 states that the probability our bound holds for one fixed pair (k, ℓ) is at least

$$1 - C(n)e^{-\Omega(n/C(n)^{1+2\alpha})} \geq 1 - O(n)e^{-\Omega(n^u)}$$

for some positive u . This is because $C(n) \leq B(n)n^{\varepsilon/2} \leq O(n^{1-\varepsilon/2})$ for large n , and so our choice of α ensures that $C(n)^{1+2\alpha} = O(n^{1-u})$ for some $u > 0$.

We can then perform a union bound over all pairs $k\ell \leq B(n)$: since the number of such pairs can be bounded by a polynomial in n , we have that the desired condition is satisfied for all $k\ell \leq B(n)$ with probability approaching 1. \blacksquare

We are now ready to prove the main result.

Theorem *Let $f : m^\infty \rightarrow \mathbb{R}$ be uniformly continuous, and assume that $f \circ r_i \rightarrow f$ pointwise. Then for every random variable (X, Y) , we have*

$$(f \circ r_{B(n)}) \left(\widehat{M}(D_n) \right) \rightarrow f(M(X, Y))$$

in probability where D_n is a sample of size n from the distribution of (X, Y) , provided $\omega(1) < B(n) \leq O(n^{1-\varepsilon})$ for some $\varepsilon > 0$.

Proof Let N denote $B(n)$, let $M_N = r_N(M)$, and let $\widehat{M}_N(D_n) = r_N(\widehat{M}(D_n))$. We begin by writing

$$\begin{aligned} \left| f \left(\widehat{M}_N(D_n) \right) - f(M) \right| &\leq \left| f \left(\widehat{M}_N(D_n) \right) - f(M_N) \right| + |f(M_N) - f(M)| \\ &= \left| f \left(\widehat{M}_N(D_n) \right) - f(M_N) \right| + |(f \circ r_N)(M) - f(M)| \end{aligned}$$

and observing that as $n \rightarrow \infty$, the second term vanishes by the pointwise convergence of $f \circ r_i$ and the fact that $B(n) > \omega(1)$. It therefore suffices to show that the first term converges to zero in probability. Since f is uniformly continuous, we can establish this via a simple adaptation of the continuous mapping theorem, which says that if the sequence of random variables $R_n \rightarrow R$ in probability, and g is continuous, then $g(R_n) \rightarrow g(R)$ in probability. We replace R with a second sequence, and replace continuity with uniform continuity.

Let $\|\cdot\|$ denote the supremum norm on m^∞ , and fix any $z > 0$. Then, for any $\delta > 0$, define

$$C_\delta = \{A \in m^\infty : \exists A' \in m^\infty \text{ s.t. } \|A - A'\| < \delta, |f(A) - f(A')| > z\}.$$

This is the set of matrices $A \in m^\infty$ for which it is possible to find, within a δ -neighborhood of A , a second matrix that f maps to more than z away from $f(A)$. Because f is uniformly continuous, there exists a δ^* sufficiently small so that $C_{\delta^*} = \emptyset$.

Suppose that $|f(\widehat{M}_N(D_n)) - f(M_N)| > z$. This means that either $\|\widehat{M}_N(D_n) - M_N\| > \delta^*$, or $M_N \in C_{\delta^*}$. The latter option is impossible since $C_{\delta^*} = \emptyset$, and Lemma 38 tells us that $\mathbf{P} \left(\|\widehat{M}_N(D_n) - M_N\| > \delta^* \right) \rightarrow 0$ as n grows. We therefore have that

$$\left| f \left(\widehat{M}_N(D_n) \right) - f(M_N) \right| \rightarrow 0$$

in probability, as desired. \blacksquare

Appendix B. Proof of Theorem 8

In this appendix we prove Theorem 8, reproduced below.

Theorem *Let $\mathcal{P}(\mathbb{R}^2)$ denote the space of random variables supported on \mathbb{R}^2 equipped with the metric of statistical distance. The map from $\mathcal{P}(\mathbb{R}^2)$ to m^∞ defined by $(X, Y) \mapsto M(X, Y)$ is uniformly continuous.*

The proposition below begins our argument with the simple observation that the family of maps consisting of applying any finite grid to some $(X, Y) \in \mathcal{P}(\mathbb{R}^2)$ is uniformly equicontinuous. The reason this holds is that $(X, Y)|_G$ is a deterministic function of (X, Y) , and deterministic functions cannot increase statistical distance.

Proposition 39 *Let \mathbb{G} be the set of all finite grids. The family $\{(X, Y) \mapsto (X, Y)|_G : G \in \mathbb{G}\}$ is uniformly equicontinuous on $\mathcal{P}(\mathbb{R}^2)$.*

Proof To establish uniform equicontinuity, we need to show that, given some $(X, Y) \in \mathcal{P}(\mathbb{R}^2)$ and some $\varepsilon > 0$, we can choose δ to satisfy the continuity condition in a way that does not depend on G or on (X, Y) . But because deterministic functions cannot increase statistical distance, we have that if $(X, Y), (X', Y') \in \mathcal{P}$ are at most ε apart then

$$\Delta((X, Y)|_G, (X', Y')|_G) \leq \Delta((X, Y), (X', Y')) = \varepsilon$$

where Δ denotes statistical distance. Choosing $\delta = \varepsilon$ therefore gives the result. ■

At this point it is tempting to try to use continuity properties of discrete mutual information to obtain uniform continuity of the characteristic matrix. And indeed, this strategy does yield that each *individual* entry of the characteristic matrix is a uniformly continuous function. However, to obtain continuity of the entire (infinite) characteristic matrix we need to make a statement about all grid resolutions simultaneously. This is not straightforward because mutual information is only uniformly continuous for a fixed grid resolution, and the family $\{(X, Y) \mapsto I((X, Y)|_G) : G \in \mathbb{G}\}$ is in fact not even equicontinuous.

The normalization in the definition of MIC_* is what allows us to establish the uniform continuity of the characteristic matrix despite this problem. To see why, suppose we have a distribution over a k -by- ℓ grid and we are allowed to move at most δ away in statistical distance for some small δ . The largest change in discrete mutual information that this can cause indeed increases as we increase k and ℓ . However, it turns out that we can bound the extent of this “non-uniformity”: the proposition below shows that as we move away from a distribution, the discrete mutual information can change only proportionally to the amount of mass we move, with the proportionality constant bounded by $\log \min\{k, \ell\}$. Because $\log \min\{k, \ell\}$ is the quantity by which we regularize the entries of the characteristic matrix, this is exactly enough to make the normalized matrix continuous. This proposition is the technical heart of our continuity result. And as we show in Corollary 11 when we demonstrate the non-continuity of the non-normalized characteristic matrix mutual information, our bound is tight.

Proposition 40 *Let $I_{k,\ell} : \mathcal{P}(\{1, \dots, k\} \times \{1, \dots, \ell\}) \rightarrow \mathbb{R}$ denote the discrete mutual information function on k -by- ℓ grids. For $0 < \delta \leq 1/4$, the maximal change in $I_{k,\ell}$ over any subset of $\mathcal{P}(\{1, \dots, k\} \times \{1, \dots, \ell\})$ of diameter δ (in statistical distance) is*

$$O\left(\delta \log\left(\frac{1}{\delta}\right) + \delta \log \min\{k, \ell\}\right).$$

Proof Without loss of generality, assume $k \leq \ell$, so that $\log \min\{k, \ell\} = \log k$. Let (X, Y) and (X', Y') be two random variables distributed over $\{1, \dots, k\} \times \{1, \dots, \ell\}$ that are at most δ apart in statistical distance. Using $I(X, Y) = H(Y) - H(Y|X)$, we can express the difference between the mutual information of these two pairs of random variables as

$$|I(X, Y) - I(X', Y')| \leq |H(Y) - H(Y')| + |H(Y|X) - H(Y'|X')|.$$

We now use Lemma 55, which relates movement of probability mass to changes in entropy and is proven in Appendix L, to separately bound each of the terms on the right hand side. Straightforward application of the lemma to $|H(Y) - H(Y')|$ shows that it is at most $2H_b(2\delta) + 3\delta \log k$, where $H_b(\cdot)$ is the binary entropy function. Since $H_b(x) \leq O(x \log(1/x))$ for x small, this is $O(\delta \log(1/\delta) + \delta \log k)$.

Bounding the term with the conditional entropies is more involved. Let $p_x = \mathbf{P}(X = x)$, and let $p'_x = \mathbf{P}(X' = x)$. We have

$$\begin{aligned} |H(Y|X) - H(Y'|X')| &= \sum_x |p_x H(Y|X = x) - p'_x H(Y'|X' = x)| \\ &\leq \sum_x (p_x |H(Y|X = x) - H(Y'|X' = x)| + |p'_x - p_x| H(Y'|X' = x)) \\ &= \sum_x p_x |H(Y|X = x) - H(Y'|X' = x)| + \sum_x |p'_x - p_x| \log k \\ &\leq \sum_x p_x |H(Y|X = x) - H(Y'|X' = x)| + \delta \log k \end{aligned} \tag{1}$$

where the last line is because $\sum_x |p_x - p'_x| \leq \delta$ and $H(Y'|X' = x) \leq \log k$.

Now let δ_{x+} be the magnitude of all the probability mass entering any cell in column x , let δ_{x-} be the magnitude of all the probability mass leaving any cell in column x , and let $\delta_x = \delta_{x+} + \delta_{x-}$. Using this notation, we can again apply Lemma 55 to obtain

$$\begin{aligned} \sum_x p_x |H(Y|X = x) - H(Y'|X' = x)| &\leq \sum_x p_x \left(2H_b\left(\frac{2\delta_x}{p_x}\right) + 3\frac{\delta_x}{p_x} \log k\right) \\ &= 2 \sum_x p_x H_b\left(\frac{2\delta_x}{p_x}\right) + 3 \sum_x \delta_x \log k \\ &\leq 2 \sum_x p_x H_b\left(\frac{2\delta_x}{p_x}\right) + 3\delta \log k \\ &\leq 2H_b(2\delta) + 3\delta \log k \end{aligned}$$

where the last line is by application of Lemma 52 from the appendix, which bounds weighted sums of binary entropies.

Combining this with Line (2) gives that

$$|H(Y|X) - H(Y'|X')| \leq 2H_b(2\delta) + 4\delta \log k$$

which, together with the bound on $|H(Y) - H(Y')|$ and the fact that $H_b(x) \leq O(x \log(1/x))$ for x small, gives the result. \blacksquare

Having bounded the extent to which variation in mutual information depends on grid resolution, we are now ready to show the uniform continuity of the characteristic matrix.

Theorem *Let $\mathcal{P}(\mathbb{R}^2)$ denote the space of random variables supported on \mathbb{R}^2 equipped with the metric of statistical distance. The map from $\mathcal{P}(\mathbb{R}^2)$ to m^∞ defined by $(X, Y) \mapsto M(X, Y)$ is uniformly continuous.*

Proof We complete the proof in three steps. First, we show that a certain family of functions F is uniformly equicontinuous. Second, we use this to show that a different family F' consisting of functions of the form $\sup_{g \in A} g$ with $A \subset F$ is uniformly equicontinuous. Finally, we argue that since the entries of $M(X, Y)$ consist of the functions in F' , this is sufficient to establish the result.

Define

$$F = \left\{ (X, Y) \mapsto \frac{I_{k,\ell}((X, Y)|_G)}{\log \min\{k, \ell\}} : k, \ell \in \mathbb{Z}_{>1}, G \in \mathcal{G}(k, \ell) \right\}.$$

F is uniformly equicontinuous by the following argument. Given some $\varepsilon > 0$, we know (Proposition 39) that for any (X', Y') in an ε -ball around (X, Y) , $(X', Y')|_G$ will remain within ε of $(X, Y)|_G$ for any G . Proposition 40 then tells us that if ε is sufficiently small then the distance between $I_{k,\ell}((X', Y')|_G)$ and $I_{k,\ell}((X, Y)|_G)$ will be at most

$$O(\varepsilon \log(1/\varepsilon) + \varepsilon \log \min\{k, \ell\}).$$

After the normalization, this becomes at most $O(\varepsilon(\log(1/\varepsilon) + 1))$, which goes to zero (uniformly with respect to (X, Y)) as ε approaches zero, as desired.

Next, define

$$F' = \{(X, Y) \mapsto M(X, Y)_{k,\ell} : k, \ell \in \mathbb{Z}_{>1}\}.$$

Each map in F' is of the form $\sup_{g \in A} g$ for some $A \subset F$. Therefore, for a given $\varepsilon > 0$, whatever δ establishes the uniform equicontinuity for F can be used to establish continuity of all the functions in F' . (To see this: $\sup_{g \in A} g$ can't increase by more than ε if no g increases by more than ε , and $\sup_{g \in A} g$ is also lower bounded by any of the g 's, so it can't decrease by more than ε either.) Since we can use the same δ for all of the maps in F' , they therefore form a uniformly equicontinuous family.

Finally, the δ provided by the uniform equicontinuity of F' also ensures that $M(X', Y')$ is within ε of $M(X, Y)$ in the supremum norm, thus giving the uniform continuity of $(X, Y) \mapsto M(X, Y)$. \blacksquare

Appendix C. Proof of Proposition 10

Theorem For some function $N(k, \ell)$, let M^N be the characteristic matrix with normalization N , i.e.,

$$M^N(X, Y) = \frac{I^*((X, Y), k, \ell)}{N(k, \ell)}.$$

If $N(k, \ell) = o(\log \min\{k, \ell\})$ along some infinite path in $\mathbb{N} \times \mathbb{N}$, then M^N and $\sup M^N$ are not continuous as functions of $\mathcal{P}([0, 1] \times [0, 1]) \subset \mathcal{P}(\mathbb{R}^2)$.

Proof Consider a random variable Z uniformly distributed on $[0, 1/2]^2$. Because Z exhibits statistical independence, $I^*(Z, k, \ell)$ is zero for all k, ℓ . Now define Z_ε to be uniformly distributed on $[0, 1/2]^2$ with probability $1 - \varepsilon$ and uniformly distributed on the line from $(1/2, 1/2)$ to $(1, 1)$ with probability ε .

We lower-bound $I^*(Z_\varepsilon, k, \ell)$. Without loss of generality suppose that $k \leq \ell$, and consider a grid that places all of $[0, 1/2]^2$ into one cell and uniformly partitions the set $[1/2, 1]^2$ into $k - 1$ rows and $k - 1$ columns. By considering just the rows/columns in the set $[1/2, 1]^2$ we see that this grid gives a mutual information of at least $\varepsilon \log(k - 1)$. Thus, we have that for all k, ℓ ,

$$I^*(Z_\varepsilon, k, \ell) \geq \varepsilon \log \min\{k - 1, \ell - 1\}.$$

This implies that the limit of $M^N(Z_\varepsilon)$ along P is ∞ , and so the distance between $M^N(Z)$ and $M^N(Z_\varepsilon)$ in the supremum norm is infinite. \blacksquare

Appendix D. Proof of Theorem 16

Theorem Let M be a population characteristic matrix. Then $M_{k,\uparrow}$ equals

$$\max_{P \in \mathcal{P}(k)} \frac{I(X, Y|_P)}{\log k}$$

where $\mathcal{P}(k)$ denotes the set of all partitions of size at most k .

Proof Define

$$M_{k,\uparrow}^* = \max_{P \in \mathcal{P}(k)} \frac{I(X, Y|_P)}{\log k}.$$

We wish to show that $M_{k,\uparrow}^*$ is in fact equal to $M_{k,\uparrow}$. To show that $M_{k,\uparrow} \leq M_{k,\uparrow}^*$, we observe that for every k -by- ℓ grid $G = (P, Q)$, where P is a partition into rows and Q is a partition into columns, the data processing inequality gives $I((X, Y)|_G) \leq I(X, Y|_P)$. Thus $M_{k,\ell} \leq M_{k,\uparrow}^*$ for $\ell \geq k$, implying that

$$M_{k,\uparrow} = \lim_{\ell \rightarrow \infty} M_{k,\ell} \leq M_{k,\uparrow}^*.$$

It remains to show that $M_{k,\uparrow}^* \leq M_{k,\uparrow}$. To do this, we let P be any partition into k rows, and we define Q_ℓ to be an equipartition into ℓ columns. We let

$$M_{k,\ell,P}^* = \frac{I(X|_{Q_\ell}, Y|_P)}{\log k}.$$

Since $M_{k,\ell,P}^* \leq M_{k,\ell}$ when $\ell \geq k$, we have that for all P

$$\frac{I(X, Y|_P)}{\log k} = \lim_{\ell \rightarrow \infty} M_{k,\ell,P}^* \leq \lim_{\ell \rightarrow \infty} M_{k,\ell} = M_{k,\uparrow}$$

which gives that

$$M_{k,\uparrow}^* = \sup_P \frac{I(X, Y|_P)}{\log k} \leq M_{k,\uparrow}$$

as desired. ■

Appendix E. Proof of Theorem 18

Theorem *Given a random variable (X, Y) , $M_{k,\uparrow}$ (resp. $M_{\uparrow,\ell}$) is computable to within an additive error of $O(k\varepsilon \log(1/(k\varepsilon))) + E$ (resp. $O(\ell\varepsilon \log(1/(\ell\varepsilon))) + E$) in time $O(kT(E)/\varepsilon)$ (resp. $O(\ell T(E)/\varepsilon)$), where $T(E)$ is the time required to numerically compute the mutual information of a continuous distribution to within an additive error of E .*

Proof Without loss of generality we prove the claim only for $M_{k,\uparrow}$. Given $0 < \varepsilon < 1$, we would like a partition into rows P of size at most k such that $I(X, Y|_P)$ is maximized. We would like to use OPTIMIZEXAXIS for this purpose, but while our search problem is continuous, OPTIMIZEXAXIS can only perform a discrete search over sub-partitions of some master partition Π . We therefore set Π to be an equipartition into $1/\varepsilon$ rows and show that this gets us close enough to achieve the desired result.

With Π as described, the OPTIMIZEXAXIS provides in time $O(kT(E)/\varepsilon)$ a partition P_0 into at most k rows such that $I(X, Y|_{P_0})$ is maximized, subject to $P_0 \subset \Pi$, to within an additive error of E . To prove the claim then, we must show that the loss we incur by restricting to sub-partitions of Π costs us at most $O(k\varepsilon \log(1/(k\varepsilon)))$. In other words, we must show that

$$I(X, Y|_P) - I(X, Y|_{P_0}) \leq O(k\varepsilon)$$

where P is an optimal partition into rows. Note that we have omitted the absolute value above, since by the optimality of P , $I(X, Y|_P) \geq I(X, Y|_{P_0})$ always.

We prove the desired bound by showing that there exists some $P' \subset \Pi$ such that the mutual information of $(X, Y|_{P'})$ is $O(k\varepsilon \log(1/(k\varepsilon)))$ -close to that achieved with $(X, Y|_P)$. Since $P' \subset \Pi$ gives us that $I(X, Y|_{P_0}) \geq I(X, Y|_{P'})$, we may then conclude that $I(X, Y|_P) - I(X, Y|_{P_0})$ is at most $O(k\varepsilon \log(1/(k\varepsilon)))$.

We construct P' by simply replacing every horizontal line in P with a horizontal line in Π closest to it. Since there are at most $k - 1$ horizontal lines in P , and each such line is contained in a row of Π containing $1/\varepsilon$ probability mass, performing this operation moves at most $(k - 1)\varepsilon$ probability mass. In other words, the statistical distance between $(X, Y|_{P'})$ and $(X, Y|_P)$ is at most $(k - 1)\varepsilon \leq k\varepsilon$. Thus, for sufficiently small ε , Proposition 40, proven in Appendix B, can be used to show that

$$|I(X, Y|_{P'}) - I(X, Y|_P)| \leq O\left(k\varepsilon \log\left(\frac{1}{k\varepsilon}\right) + k\varepsilon \log\left(\frac{1}{\varepsilon}\right)\right)$$

which yields the desired result. ■

Remark 41 *We do not explore here the details of the numerical integration associated with the above theorem, since the error introduced by the numerical integration is independent of the algorithm being proposed. However, standard numerical integration methods can be used to make this error arbitrarily small with an understood complexity tradeoff (see, e.g., [Stoer and Bulirsch 1980](#)).*

Appendix F. Proof of Theorem 21

Theorem *Let (X, Y) be jointly distributed random variables. Then $\partial[M] = \partial M$.*

Proof Without loss of generality, we show that $[M]_{k,\uparrow} = M_{k,\uparrow}$. Fix any partition into rows P . If Q_ℓ is an equipartition into ℓ columns then

$$\lim_{\ell \rightarrow \infty} I(X|_{Q_\ell}, Y|_P) = I(X, Y|_P),$$

because the continuous mutual information equals the limit of the discrete mutual information with increasingly fine partitions. (See, e.g., Chapter 8 of [Cover and Thomas 2006](#) for a proof of this.) This means that, letting $P(k)$ denote the set of all partitions of size at most k , we have

$$[M]_{k,\uparrow} = \max_{P \in P(k)} \frac{I(X, Y|_P)}{\log k} = M_{k,\uparrow}$$

where the second equality follows from Proposition 16. ■

Appendix G. Consistency of MIC_e in Estimating MIC_*

The consistency of MIC_e for estimating MIC_* can be established using the same technical lemmas that we used to show that $\text{MIC} \rightarrow \text{MIC}_*$. Specifically, we can use Lemma 37, which bounds the difference, for all k -by- ℓ grids G , between the sample quantity $I(D_n|_G)$ and the population quantity $I((X, Y)|_G)$ with high probability, where D_n is a sample of size n from (X, Y) . That lemma yields the following fact about the sample equicharacteristic matrix, whose proof is similar to that of Lemma 38.

Lemma 42 *Let D_n be a sample of size n from the distribution of a pair (X, Y) of jointly distributed random variables. For every $B(n) = O(n^{1-\varepsilon})$, there exists an $a > 0$ such that for sufficiently large n ,*

$$\left| [\widehat{M}](D_n)_{k,\ell} - [M](X, Y)_{k,\ell} \right| \leq O\left(\frac{1}{n^a}\right)$$

holds for all $k\ell \leq B(n)$ with probability $P(n) = 1 - o(1)$, where $[\widehat{M}](D_n)_{k,\ell}$ is the k, ℓ -th entry of the sample equicharacteristic matrix and $[M](X, Y)_{k,\ell}$ is the k, ℓ -th entry of the population equicharacteristic matrix of (X, Y) .

In the case of MIC, we proceeded to apply abstract continuity considerations to obtain our consistency theorem (Theorem 6) from a result analogous to the above lemma. A similar argument shows us that, in the case of the equicharacteristic matrix as well, we can estimate a large class of functions of the matrix in the same way. This is stated formally in the theorem below. As before, we let m^∞ be the space of infinite matrices equipped with the supremum norm, and given a matrix A the projection r_i zeros out all the entries $A_{k,\ell}$ for which $k\ell > i$.

Theorem *Let $f : m^\infty \rightarrow \mathbb{R}$ be uniformly continuous, and assume that $f \circ r_i \rightarrow f$ pointwise. Then for every random variable (X, Y) , we have*

$$(f \circ r_{B(n)}) \left([\widehat{M}](D_n) \right) \rightarrow f([M](X, Y))$$

in probability where D_n is a sample of size n from the distribution of (X, Y) , provided $\omega(1) < B(n) \leq O(n^{1-\varepsilon})$ for some $\varepsilon > 0$.

Appendix H. The EQUICHARCLUMP Algorithm

In Theorem 28, we sketched an algorithm called EQUICHARCLUMP for approximating the sample equicharacteristic matrix that is more efficient than the naive computation. In this appendix, we describe the algorithm in detail, bound its runtime, and show that it indeed yields a consistent estimator of MIC_* from finite samples as well as a consistent independence test when used to compute the total information coefficient. We then present some empirical results characterizing the sensitivity of the algorithm to its speed-versus-optimality parameter c .

The results in this section can be summarized as follows: let (X, Y) be a pair of jointly distributed random variables, and let D_n be a sample of size n from the distribution of (X, Y) . For every $c \geq 1$, there exists a matrix $\{\widehat{M}\}^c(D_n)$ such that

1. There exists an algorithm EQUICHARCLUMP for computing $r_B(\{\widehat{M}\}^c(D_n))$ in time $O(n + B^{5/2})$, which equals $O(n + n^{5(1-\varepsilon)/2})$ when $B(n) = O(n^{1-\varepsilon})$.

2. The function

$$\widetilde{\text{MIC}}_{e,B}(\cdot) = \max_{k\ell \leq B(n)} \{\widehat{M}\}^c(\cdot)_{k,\ell}$$

is a consistent estimator of MIC_* provided $\omega(1) < B(n) \leq O(n^{1-\varepsilon})$ for some $\varepsilon > 0$.

3. The function

$$\widetilde{\text{TIC}}_{e,B}(\cdot) = \sum_{k\ell \leq B(n)} \{\widehat{M}\}^c(\cdot)_{k,\ell}$$

yields a consistent right-tailed test of independence provided $\omega(1) < B(n) \leq O(n^{1-\varepsilon})$ for some $\varepsilon > 0$

We will prove these results in order.

H.1 Algorithm Description and Analysis of Runtime

We begin by describing the algorithm and bounding its runtime simultaneously. As in the proof of Theorem 27, we bound the runtime required to approximately compute only the k, ℓ -th entries of $\{\widehat{M}\}^c(D_n)$ satisfying $k \leq \ell, k\ell \leq B$. To do this, we analyze two portions of $\{\widehat{M}\}^c(D_n)$ separately: we first consider the case $\ell \geq \sqrt{B}$, in which we must compute the entries corresponding to all the pairs $\{(2, \ell), \dots, (B/\ell, \ell)\}$. We then consider $\ell < \sqrt{B}$, in which case we need only compute the entries $\{(2, \ell), \dots, (\ell, \ell)\}$ since the additional pairs would all have $k > \ell$.

For the case of $\ell \geq \sqrt{B}$, as in the previous theorem we can simultaneously compute using OPTIMIZEXAXIS the entries corresponding to all the pairs $\{(2, \ell), \dots, (B/\ell, \ell)\}$ in time $O(|\Pi|^2(B/\ell)\ell) = O(|\Pi|^2B)$, which equals $O(c^2B^3/\ell^2)$ when we set Π to be an equipartition of size cB/ℓ . Doing this for $\ell = \sqrt{B}, \dots, B/2$ gives a contribution of the following order to the runtime.

$$\begin{aligned} O(c^2B^3) \sum_{\ell=\sqrt{B}}^{B/2} \frac{1}{\ell^2} &= O(c^2B^3) O\left(\frac{1}{\sqrt{B}}\right) \\ &= O(c^2B^{5/2}) \end{aligned}$$

For the case of $\ell < \sqrt{B}$, we can simultaneously compute using OPTIMIZEXAXIS the entries corresponding to all the pairs $\{(2, \ell), \dots, (\ell, \ell)\}$ in time $O(|\Pi|^2\ell^2)$ which equals $O(c^2\ell^4) \leq O(c^2B^2)$ when we set Π to be an equipartition of size $c\ell$. Summing over the $O(\sqrt{B})$ possible values of ℓ with $\ell < \sqrt{B}$ gives an upper bound of $O(c^2B^{5/2})$.

H.2 Consistency

Let (X, Y) be a pair of jointly distributed random variables. For a sample D_n of size n from the distribution of (X, Y) and a speed-versus-optimality parameter $c \geq 1$, let $\{\widehat{M}\}^c(D_n)$ denote the matrix computed by EQUICHARCLUMP. (Notice the use of curly braces to differentiate this from the sample equicharacteristic matrix $[\widehat{M}]$.) We show here that $\max_{k\ell \leq B(n)} \{\widehat{M}\}^c(D_n)_{k,\ell}$ is a consistent estimator of $\text{MIC}_*(X, Y)$, and correspondingly that $\sum_{k\ell \leq B(n)} \{\widehat{M}\}^c(D_n)_{k,\ell}$ yields a consistent independence test.

The key to both consistency results is that, though in calculating the k, ℓ -th entry of $\{\widehat{M}\}^c(D_n)$ the algorithm only searches for optimal partitions that are sub-partitions of some equipartition, the size of the equipartition used always grows as n, k , and ℓ grow large. Therefore, in the limit this additional restriction does not hinder the optimization. We present this argument by introducing a population object called the *clumped equicharacteristic matrix*. We observe that this matrix is the limit of the EQUICHARCLUMP procedure as sample size grows, and then show that the supremum and partial sums of this matrix have the necessary properties.

Definition 43 *Let (X, Y) be jointly distributed random variables and fix some $c \geq 1$. Let*

$$I^{\{c^*\}}((X, Y), k, \ell) = \max_G I((X, Y)|_G)$$

where the maximum is over k -by- ℓ grids whose larger partition is an equipartition and whose smaller partition must be contained in an equipartition of size $c \cdot \max\{k, \ell\}$. The clumped equicharacteristic matrix of (X, Y) , denoted by $\{M\}^c(X, Y)$, is defined by

$$\{M\}^c(X, Y)_{k,\ell} = \frac{I^{\{c^*\}}((X, Y), k, \ell)}{\log \min\{k, \ell\}}$$

Notice that curly braces differentiate the quantities $I^{\{c^*\}}$ and $\{M\}^c$ defined above from the corresponding equicharacteristic matrix quantities $I^{[*]}$ and $[M]$.

The following two results, which we state without proof, characterize the convergence of the output of EQUICHARCLUMP to the clumped equicharacteristic matrix. These lemmas can be shown using Lemma 37, which simultaneously bounds the difference, for all k -by- ℓ grids G , between the sample quantity $I(D_n|_G)$ and the population quantity $I((X, Y)|_G)$ with high probability over the sample D_n of size n from (X, Y) .

Lemma 44 *Let D_n be a sample of size n from the distribution of a pair (X, Y) of jointly distributed random variables. For every $B(n) = O(n^{1-\varepsilon})$, there exists an $a > 0$ such that for sufficiently large n ,*

$$\left| \{\widehat{M}\}^c(D_n)_{k,\ell} - \{M\}^c(X, Y)_{k,\ell} \right| \leq O\left(\frac{1}{n^a}\right)$$

holds for all $k, \ell \leq \sqrt{B(n)}$ with probability $P(n) = 1 - o(1)$, where $\{\widehat{M}\}^c(D_n)$ denotes the matrix computed by the EQUICHARCLUMP algorithm with parameter c on the sample D_n .

Notice that the error bound provided by the above lemma holds not for $k\ell \leq B(n)$ as in the analogous Lemma 38 and Lemma 42, but rather for the smaller region defined by $k, \ell \leq \sqrt{B(n)}$. However, though we do not have uniform convergence outside the region $k, \ell \leq \sqrt{B(n)}$, we do nevertheless have pointwise convergence there, as stated below.

Lemma 45 *Fix $k, \ell \geq 2$. Let D_n be a sample of size n from the distribution of a pair (X, Y) of jointly distributed random variables. For every $B(n) > \omega(1)$, we have that*

$$\{\widehat{M}\}^c(D_n)_{k,\ell} \rightarrow \{M\}^c(X, Y)_{k,\ell}$$

in probability as n grows, where $\{\widehat{M}\}^c(D_n)$ denotes the matrix computed by the EQUICHARCLUMP algorithm with parameter c on the sample D_n .

H.2.1 CONSISTENCY FOR ESTIMATING MIC_*

The consistency of $\{\widehat{M}\}^c(D_n)$ for estimating MIC_* follows from the following property of the clumped equicharacteristic matrix $\{M\}^c$, for which we state a proof sketch.

Proposition 46 *Let (X, Y) be a pair of jointly distributed random variables. Then we have $\sup\{M\}^c(X, Y) = MIC_*(X, Y)$.*

Proof (Sketch) Let $\{M\}^c = \{M\}^c(X, Y)$, and let $M = M(X, Y)$ be the characteristic matrix. Fix k , and consider the limit $\{M\}_{k, \ell}^c$ as ℓ grows. The grid chosen for the k, ℓ -th entry when $\ell > k$ will contain an equipartition P_ℓ of size ℓ on the x-axis, and a partition Q_ℓ of size k on the y-axis that is optimal subject to the restriction that Q_ℓ be contained in an equipartition of size $c\ell$. As ℓ grows large, the equipartition P_ℓ on the first axis will become finer and finer until in the limit $X|_{P_\ell} \rightarrow X$. And the partition Q_ℓ will be chosen from a finer and finer equipartition, so that in the limit it approaches an unconditionally optimal partition Q of size k . The convergence of Q_ℓ to the optimal partition Q of size k can be shown to be uniform using Proposition 40. This implies that

$$\{M\}_{k, \uparrow}^c = \lim_{\ell \rightarrow \infty} \{M\}_{k, \ell}^c = \max_{P \in P(k)} \frac{I(X, Y|_P)}{\log k}$$

where $P(k)$ denotes the set of all partitions of size at most k . Therefore, the boundary $\partial\{M\}^c$ of $\{M\}^c$ equals the boundary ∂M of M . Since $\text{MIC}_*(X, Y) = \sup \partial M$ (Theorem 15), this implies that

$$\sup\{M\}^c \geq \sup \partial\{M\}^c = \sup \partial M = \text{MIC}_*(X, Y).$$

On the other hand, $\{M\}^c \leq M$ element-wise since the optimization for the k, ℓ -th entry of $\{M\}^c$ is performed over a subset of the grids searched for the k, ℓ -th entry of M . This means that $\sup\{M\}^c \leq \sup M = \text{MIC}_*(X, Y)$. \blacksquare

This fact, together with the pointwise convergence of $\{\widehat{M}\}^c(D_n)$ to $\{M\}^c$, suffices to establish the consistency we seek via standard continuity arguments, which we give in the abstract lemma below. The lemma applies to a double-indexed sequence indexed by i and j ; in our argument, the index i corresponds to position in the equicharacteristic matrix, and the index j corresponds to sample size. The sequence A corresponds to the output of the EQUICHARCLUMP algorithm, the sequence a corresponds to the clumped equicharacteristic matrix, and the sequence B corresponds to the sample equicharacteristic matrix.

Lemma 47 *Let $\{A_{ij}\}_{i,j=1}^\infty$ and $\{B_{ij}\}_{i,j=1}^\infty$ be sequences of random variables, and let $\{a_i\}_{i=1}^\infty$ be a non-stochastic sequence. Assume that the following conditions hold.*

1. $A_{ij} \leq B_{ij}$ almost surely
2. For every i , $A_{ij} \rightarrow a_i$ in probability
3. $B'_j = \max_{i \leq j} B_{ij}$ satisfies $B'_j \rightarrow \sup\{a_i\}$ in probability

Then $A'_j = \max_{i \leq j} A_{ij}$ converges in probability to $\sup\{a_i\}$ as well.

Proof Let $a = \sup\{a_i\}$. We give the proof for the case that $a < \infty$. However, it is easily adapted to the infinite case. We must show that for every $\varepsilon > 0$ and every $0 < p \leq 1$, there exists some N such that $\mathbf{P}(|A'_j - a| < \varepsilon) > p$ for all $j \geq N$. By the definition of a , we know that there exists some k such that $|a_k - a| < \varepsilon/2$. Also, by the convergence of A_{kj} to a_k , there exists some m such that $\mathbf{P}(|A_{kj} - a_k| < \varepsilon/2) > 1 - p$ for all $j \geq m$. Thus, with probability at least $1 - p$, we have

$$\begin{aligned} |A_{kj} - a| &\leq |A_{kj} - a_k| + |a_k - a| \\ &\leq \varepsilon \end{aligned}$$

for all $j \geq m$.

Next, we observe that since $A'_j \geq A_{kj}$ for $j \geq k$, the above inequality implies that for $j \geq \max\{m, k\}$ we have $\mathbf{P}(A'_j > a - \varepsilon) > 1 - p$. It remains only to show that A'_j doesn't get too large, but this follows from the fact that $A'_j \leq B'_j$ and $B'_j \rightarrow a$ in probability. Specifically, we are guaranteed some $N \geq \max\{m, k\}$ such that $\mathbf{P}(B'_j < a + \varepsilon) > 1 - p$ for $j \geq N$. Since $B'_j < a + \varepsilon$ implies $A'_j < a + \varepsilon$, we have that $\mathbf{P}(|A'_j - a| < \varepsilon) > 1 - p$ for $j \geq N$, as desired. ■

Proposition 48 *The function*

$$\widehat{MIC}_{e,B}(\cdot) = \max_{k\ell \leq B(n)} \{\widehat{M}\}^c(\cdot)_{k,\ell}$$

is a consistent estimator of MIC_* provided $\omega(1) < B(n) \leq O(n^{1-\varepsilon})$ for some $\varepsilon > 0$, where $\{\widehat{M}\}^c(\cdot)$ is the output of the EQUICHARCLUMP algorithm.

Proof Let (X, Y) be a pair of jointly distributed random variables, and let D_n be a sample of size n from the distribution of (X, Y) . Let $\{(k_i, \ell_i)\}_{i=1}^\infty \subset \mathbb{Z}^+ \times \mathbb{Z}^+$ be a sequence of coordinates with the property that for every number B there exists an index $q(B)$ such that

$$\{(k_i, \ell_i) : i \leq q(B)\} = \{(k, \ell) : k\ell \leq B\}.$$

We define $B_{ij} = [\widehat{M}](D_j)_{k_i, \ell_i}$, i.e., B_{ij} is the k_i, ℓ_i -th entry of the sample characteristic matrix evaluated on a sample of size j . We analogously define $A_{ij} = \{\widehat{M}\}^c(D_j)_{k_i, \ell_i}$, and we define $a_i = \{M\}^c(X, Y)_{k_i, \ell_i}$. We observe that by Proposition 46, $\sup a_i = \sup \{M\}^c(X, Y) = MIC_*$.

It is straightforward to see that $A_{ij} \leq B_{ij}$. Additionally, Lemma 45 shows that $A_{ij} \rightarrow a_i$ in probability, and Corollary 26, which states that MIC_e is a consistent estimator of MIC_* , shows that $B'_j = \max_{i \leq j} B_{ij} \rightarrow MIC_*(X, Y)$. In the notation of the lemma, it therefore follows that $A'_j = \max_{i \leq j} A_{ij}$ converges in probability to $MIC_*(X, Y)$ as well. But this means that the sub-sequence

$$A'_{q(B(n))} = \max_{i \leq q(B(n))} \{\widehat{M}\}^c(D_{q(B(n))})_{k_i, \ell_i} = \max_{k\ell \leq B(n)} \{\widehat{M}\}^c(D_{q(B(n))})_{k,\ell}$$

converges in probability to $MIC_*(X, Y)$, which implies the result since the sequence A'_j is monotone. ■

H.2.2 CONSISTENCY FOR TOTAL INFORMATION COEFFICIENT

Similarly to the consistency argument for MIC_* , we begin by exhibiting the relevant property of the population clumped equicharacteristic matrix.

Proposition 49 *Let (X, Y) be a pair of jointly distributed random variables. If X and Y are statistically independent, then $\{M\}^c(X, Y) \equiv 0$. If not, then there exists some $a > 0$ and some integer $\ell_0 \geq 2$ such that*

$$\{M\}^c(X, Y)_{k,\ell} \geq \frac{a}{\log \min\{k, \ell\}}$$

either for all $k \geq \ell \geq \ell_0$, or for all $\ell \geq k \geq \ell_0$.

Proof (Sketch) Let $\{M\}^c = \{\widehat{M}\}^c(X, Y)$. Under independence, every entry of $\{M\}^c$ is zero since $I((X, Y)|_G) = 0$ for any grid G . For the case of dependence, the argument is identical to that given in the proof of Proposition 32. Specifically, it can be shown that there exists some index ℓ_0 , taken without loss of generality to be a column index, and some $r > 0$ such that all but finitely many of the entries in the ℓ_0 -column are at least r . It can then be shown that for large k , the entries $(k, \ell_0), (k, \ell_0 + 1), \dots, (k, k)$ have non-decreasing values of $I^{[c*]}$. This establishes the claim for $a = r \log \ell_0$. \blacksquare

We now show that the above result, together with the uniform convergence of $\{\widehat{M}\}^c(D_n)$ to $\{M\}^c(X, Y)$, implies the consistency we seek.

Proposition 50 *The function*

$$\widetilde{\text{TIC}}_{e,B}(\cdot) = \sum_{k\ell \leq B(n)} \{\widehat{M}\}^c(\cdot)_{k,\ell}$$

yields a consistent right-tailed test of independence provided $\omega(1) < B(n) \leq O(n^{1-\varepsilon})$ for some $\varepsilon > 0$, where $\{\widehat{M}\}^c(\cdot)$ is the output of the the EQUICHARCLUMP algorithm.

Proof Let (X, Y) a pair of jointly distributed random variables, and let D_n be a sample of size n from the distribution of (X, Y) . It suffices to show consistency for any deterministic monotonic function of the statistic in question. We therefore choose to analyze $\widetilde{\text{TIC}}_{e,B}(D_n) \log(B(n))/B(n)$.

For the null hypothesis in which X and Y are independent, we observe that since $\{\widehat{M}\}^c(D_n) \leq \widehat{M}(D_n)$ element-wise, $0 \leq \widetilde{\text{TIC}}_{e,B}(D_n) \leq \text{TIC}_{e,B}(D_n)$ as well. Moreover, the argument given in Appendix K, which shows that $\text{TIC}_{e,B}(D_n)/B(n)$ converges to 0 in probability under the null hypothesis, can be adapted to show that $\text{TIC}_{e,B}(D_n) \log(B(n))/B(n) \rightarrow 0$ as well. Thus, $\widetilde{\text{TIC}}_{e,B}(D_n) \log(B(n))/B(n)$ converges to zero in probability, as required.

For the case that X and Y are dependent, the proof is analogous to the argument given in Appendix K for TIC_e . The only difference is that Lemma 44, which guarantees the uniform convergence of $\{\widehat{M}\}^c(D_n)$ to $\{M\}^c(X, Y)$, applies only to the k, ℓ -th entries for which $k, \ell \leq \sqrt{B(n)}$, rather than the entries over which we are summing, which are those for which $k\ell \leq B(n)$. However, since we require only a lower bound on $\widetilde{\text{TIC}}_{e,B}(D_n)$, we may neglect these entries because

$$\widetilde{\text{TIC}}_{e,B}(D_n) = \sum_{k\ell \leq B(n)} \{\widehat{M}\}^c(D_n)_{k,\ell} \geq \sum_{k,\ell \leq \sqrt{B(n)}} \{\widehat{M}\}^c(D_n)_{k,\ell}.$$

It can then be shown, following the argument from Appendix K, that there exists some $a > 0$ depending only on B such that, with probability $1 - o(1)$,

$$\frac{\log B(n)}{B(n)} \left(\sum_{k,\ell \leq \sqrt{B(n)}} \{\widehat{M}\}^c(X, Y)_{k,\ell} - \widetilde{\text{TIC}}_{e,B}(D_n) \right) \leq O\left(\frac{\#_n \log B(n)}{B(n)n^a}\right) = O\left(\frac{\log B(n)}{n^a}\right)$$

where $\#_n = B(n)$ represents the number of pairs (k, ℓ) such that $k, \ell \leq \sqrt{B(n)}$. To obtain the result, we note that this means that

$$\frac{\log B(n)}{B(n)} \widetilde{\text{TIC}}_{e,B}(D_n) \geq \frac{\log B(n)}{B(n)} \sum_{k, \ell \leq \sqrt{B(n)}} \{\widehat{M}\}^c(X, Y)_{k, \ell} - O\left(\frac{\log B(n)}{n^\alpha}\right)$$

and then invoke Proposition 49, which implies that for large n

$$\sum_{k, \ell \leq \sqrt{B(n)}} \{M\}^c(X, Y) \geq \Omega\left(\frac{B(n)}{\log B(n)}\right).$$

■

H.3 Empirical Characterization of the Performance of EQUICHARCLUMP

The EQUICHARCLUMP algorithm has a parameter c that controls the fineness of the equipartition whose sub-partitions are searched over by the algorithm. To gain an empirical understanding of the effect of c on performance, we computed MIC_e on the set of relationships described in Section 4.4 using EQUICHARCLUMP with different values of c . For each relationship, we compared the average MIC_e across all 500 independent samples from that relationship with different values of c . We performed this analysis at sample sizes of $n = 250$ (Figure H1), $n = 500$ (Figure H2), and 5,000 (Figure H3).

We summarize our findings as follows.

- At low ($n = 250$) and medium ($n = 500$) sample sizes, using $c = 1$ introduces a downward bias for more complex relationships when $B(n) = n^{0.6}$ is used but not when $B(n) = n^{0.8}$ is used. This makes sense since the low sample size and low setting of $B(n)$ mean that the algorithm is searching over grids with relatively few cells, and so setting $c = 1$ hinders its ability to find good grids in this limited search space. This bias is almost entirely alleviated by setting $c \geq 2$.
- At high sample size ($n = 5,000$), this effect is still observable but much reduced. This makes sense since when n is large, $B(n)$ is large as well, and so the number of cells allowed in the grids being searched over is already large regardless of the exponent α used in $B(n) = n^\alpha$. Thus, there is less need for the robustness provided by searching for an optimal grid.

Appendix I. Equitability and Power Analyses from Reshef et al. (2015a)

Figure I1 contains a representative equitability analysis from Reshef et al. (2015a). Figure I2 contains power curves from Reshef et al. (2015a) for a large set of leading methods.

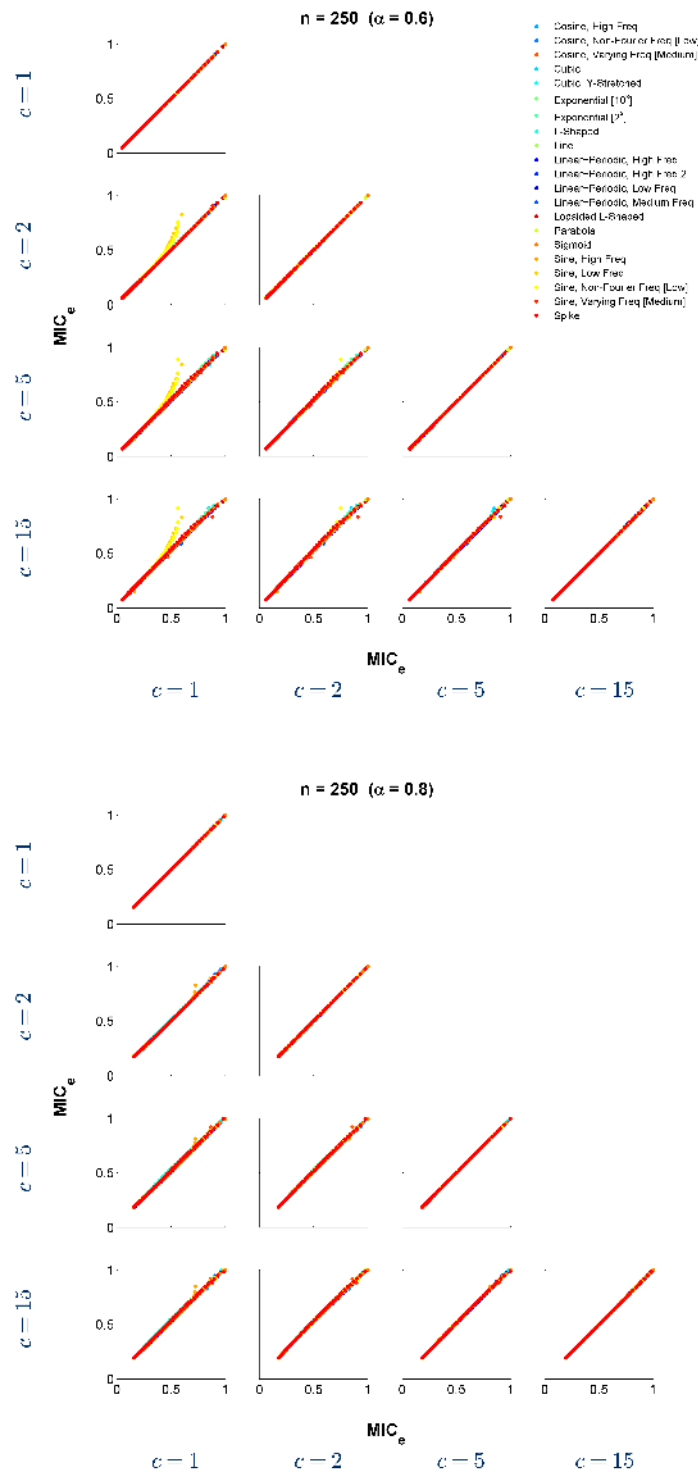


Figure H1: The effect of the parameter c on the performance of EQUICHARCLUMP, at $n = 250$. See Section H.3 for details.

MEASURING DEPENDENCE POWERFULLY AND EQUITABLY

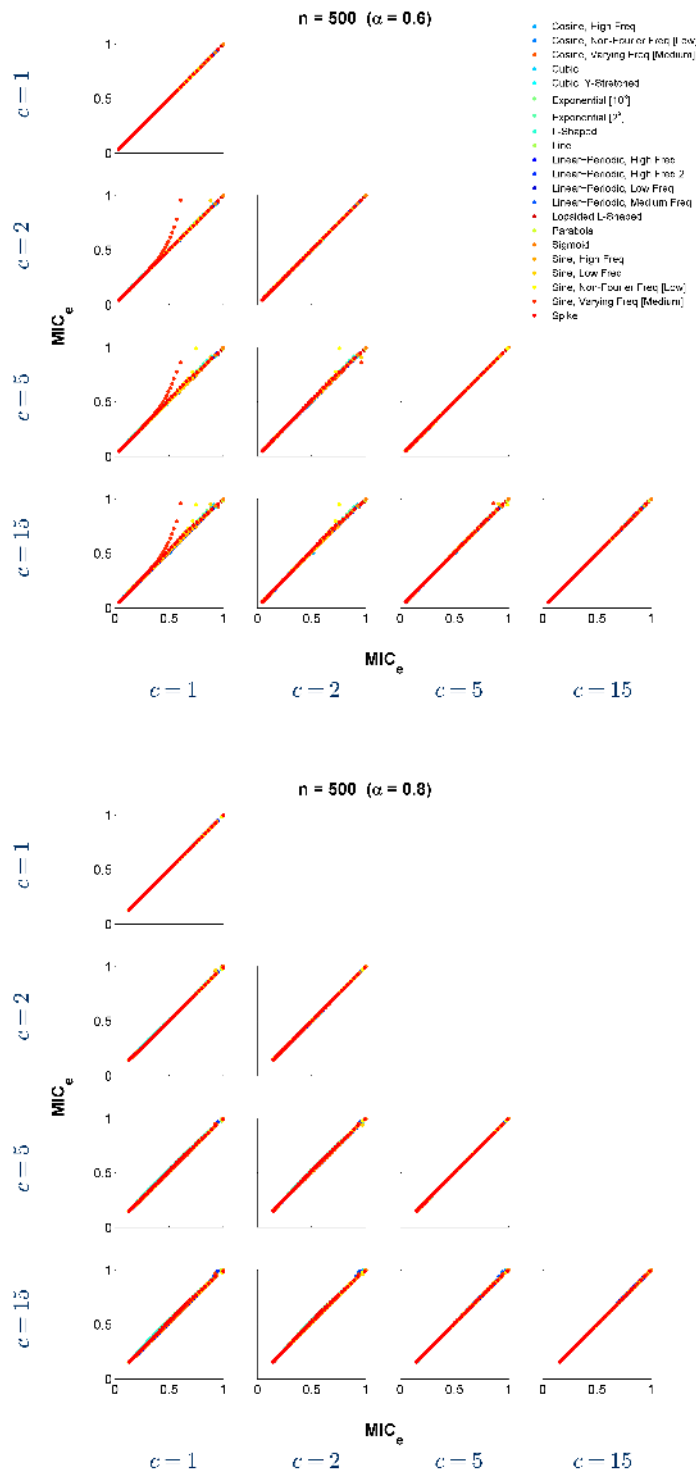


Figure H2: The effect of the parameter c on the performance of EQUICHARCLUMP, at $n = 500$. See Section H.3 for details.

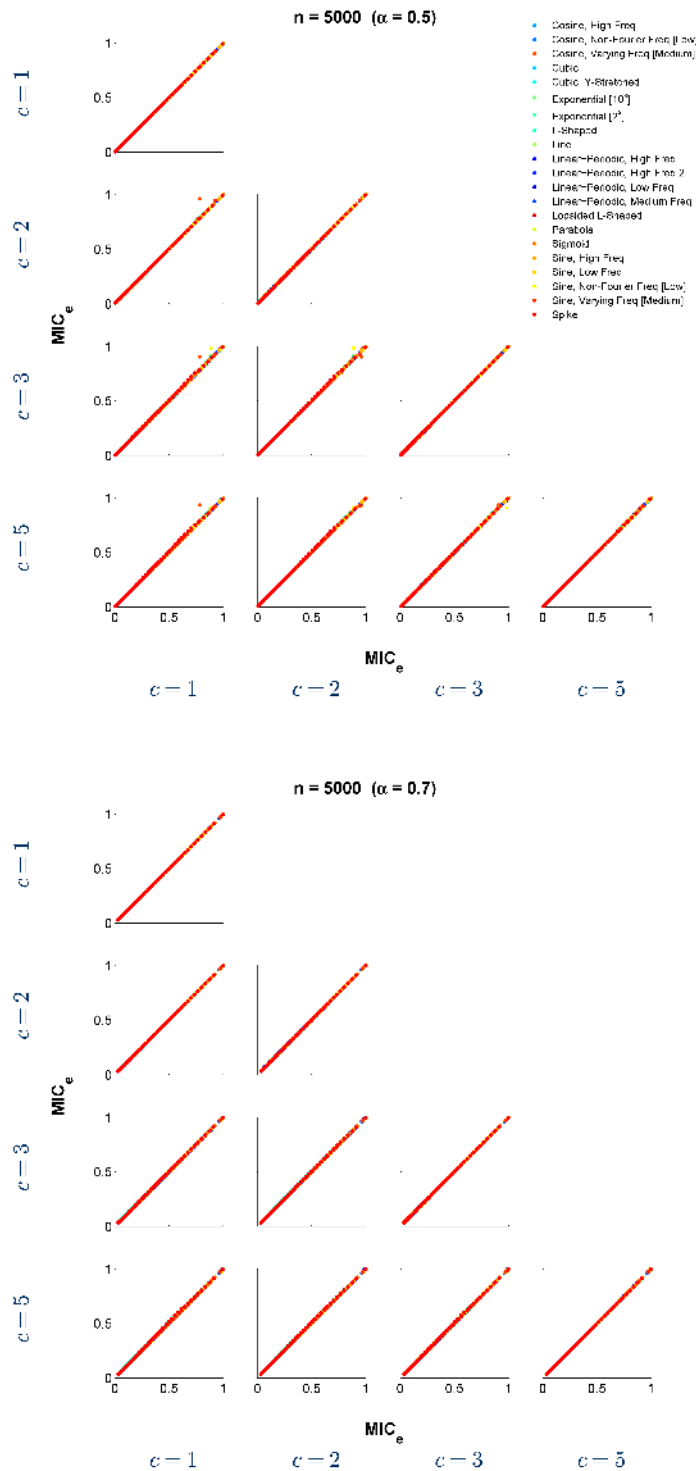


Figure H3: The effect of the parameter c on the performance of EQUICHARCLUMP, at $n = 5,000$. See Section H.3 for details.

MEASURING DEPENDENCE POWERFULLY AND EQUITABLY

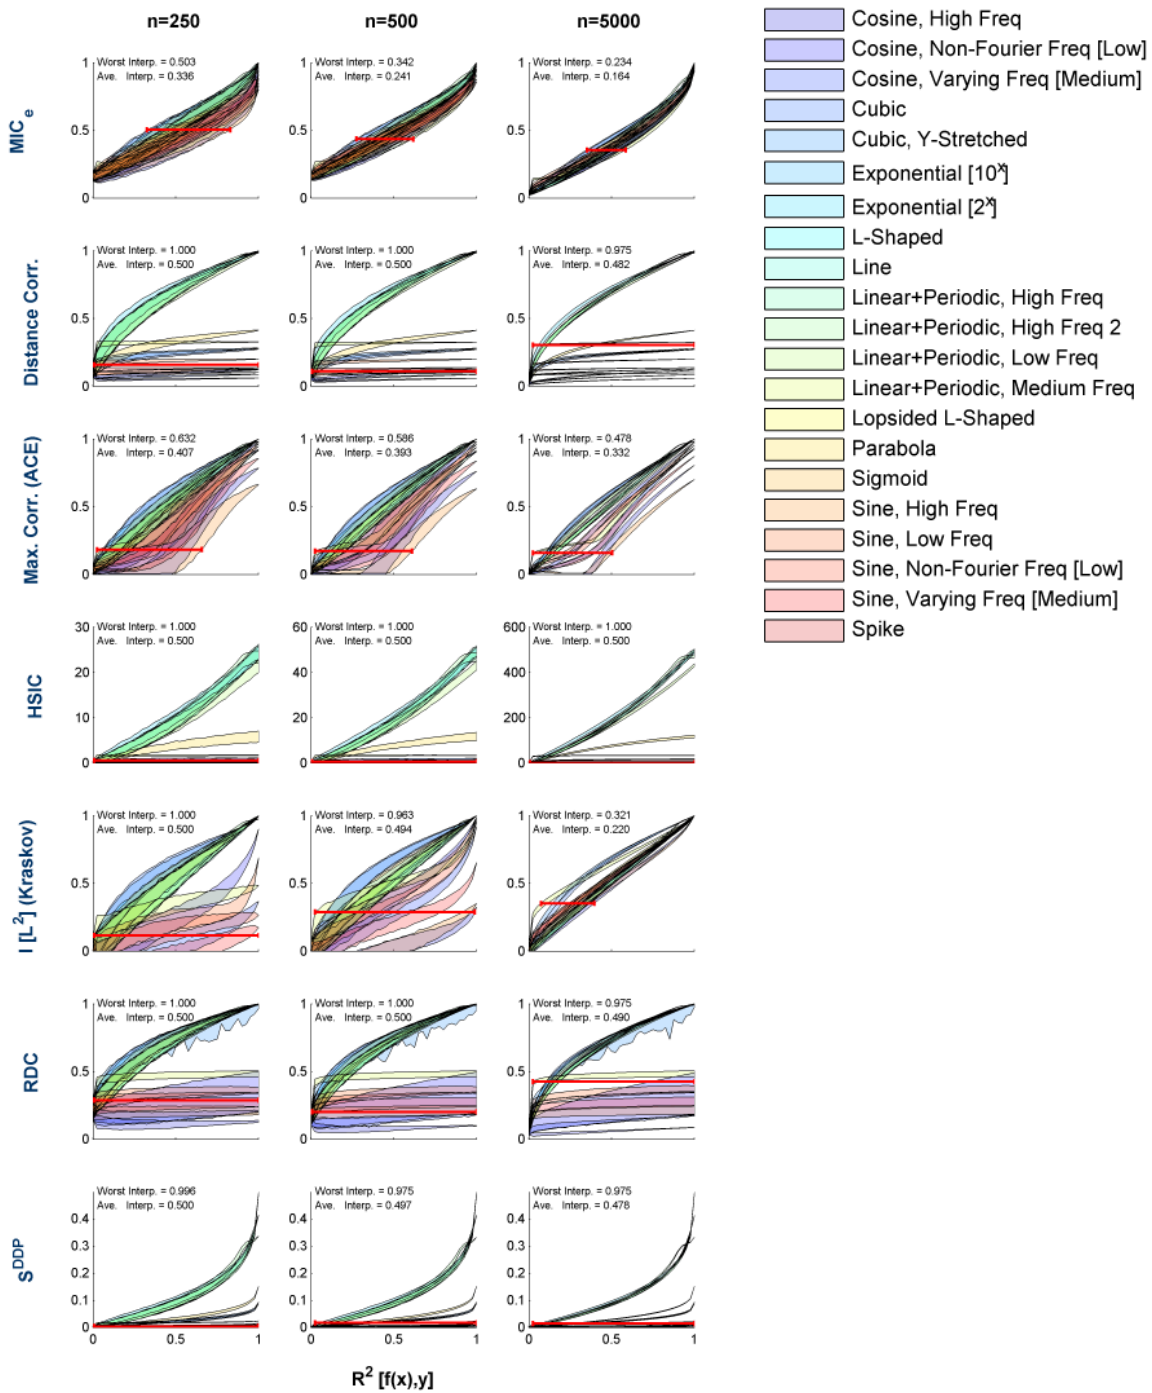


Figure I1: (Reproduced from Reshef et al., 2015a.) The equitability of measures of dependence on a set of noisy functional relationships, reproduced from Reshef et al. (2015a). [Narrower is more equitable.] The plots were constructed as in Figure 3. Mutual information, estimated using the Kraskov estimator, is represented using the squared Linfoot correlation.

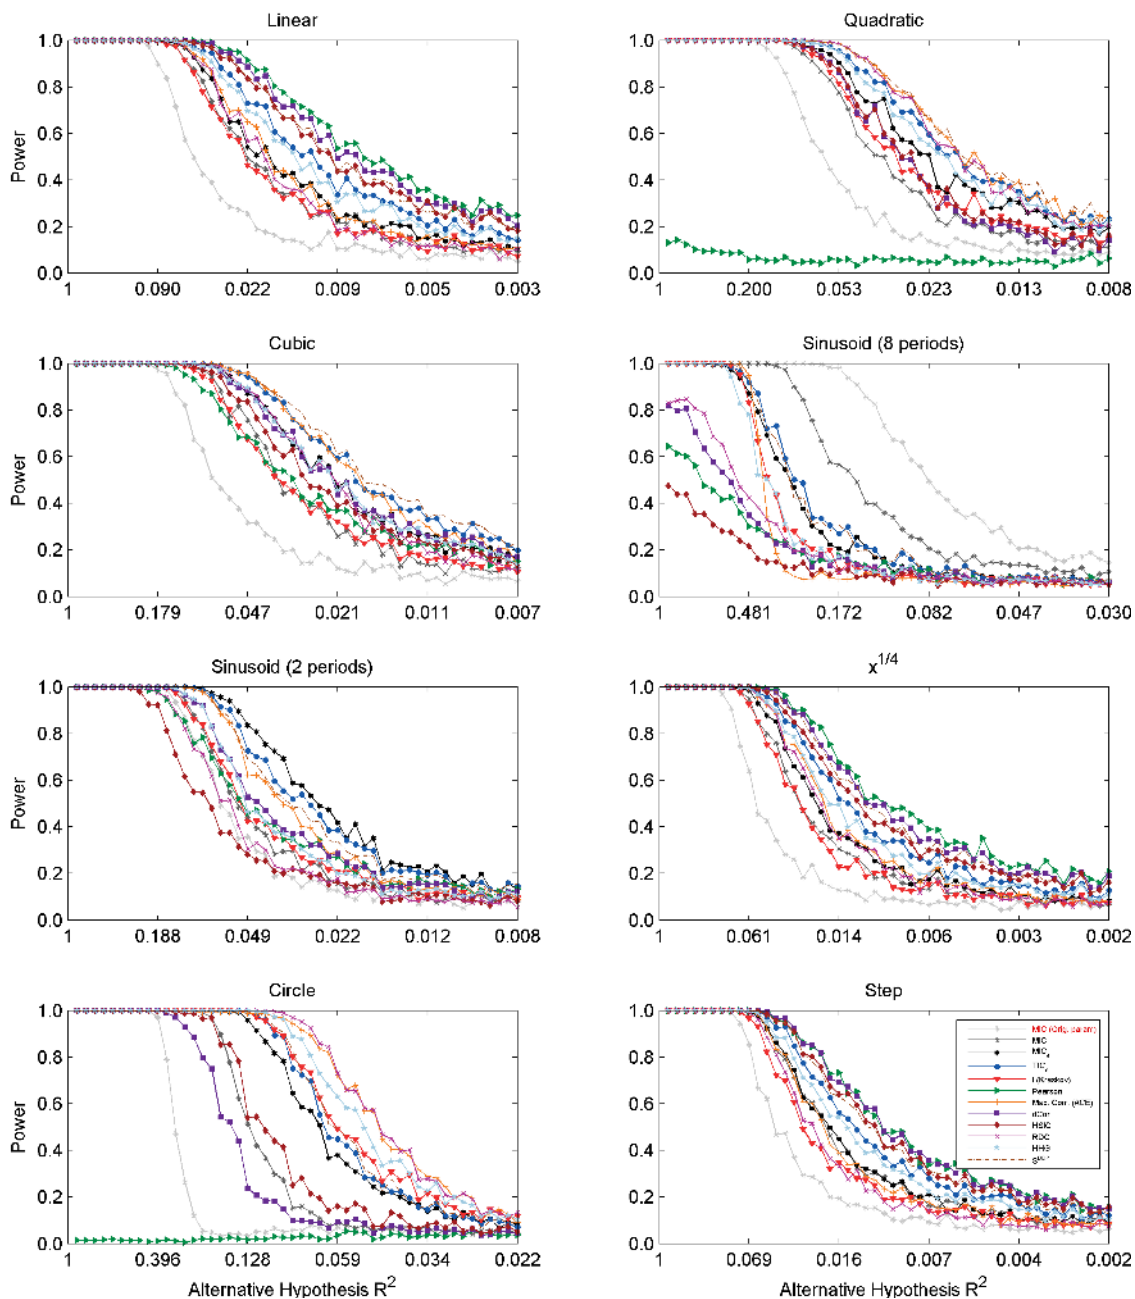


Figure I2: (Reproduced from Reshef et al., 2015a.) Power of independence testing using several leading measures of dependence, on the relationships chosen by Simon and Tibshirani (2012), at 50 noise levels with linearly increasing magnitude for each relationship and $n = 500$. To enable comparison of power regimes across relationships, the x-axis of each plot lists R^2 rather than noise magnitude.

Appendix J. Equitability Analysis of Randomly Chosen Functions with Additional Noise Model

Figure J1 contains a version of the main text Figure 4, but where noise has been added only to the dependent variable in each functional relationship, rather to both the independent and dependent variables.

Appendix K. Consistency of Independence Testing Based on TIC_e

Here we prove Propositions 32 and 33 and then use those propositions to prove Theorem 34, which shows that TIC_e can be used for independence testing.

K.1 Proof of Proposition 32

Proposition *Let (X, Y) be a pair of jointly distributed random variables. If X and Y are statistically independent, then $M(X, Y) \equiv [M](X, Y) \equiv 0$. If not, then there exists some $a > 0$ and some integer $\ell_0 \geq 2$ such that*

$$M(X, Y)_{k,\ell}, [M](X, Y)_{k,\ell} \geq \frac{a}{\log \min\{k, \ell\}}$$

either for all $k \geq \ell \geq \ell_0$, or for all $\ell \geq k \geq \ell_0$.

Proof We give the proof only for $[M] = [M](X, Y)$, with the understanding that all parts of the argument are either identical or similar for $M(X, Y)$. When X and Y are independent, then for any grid G , $(X, Y)|_G$ exhibits independence as well. Therefore $I((X, Y)|_G) = 0$ for all grids G , and so every entry of $[M]$, being a supremum over such quantities, is 0.

For the case that X and Y are dependent, our strategy is to first find, without loss of generality, a column of $[M]$ almost all of whose values are bounded away from zero, and then argue that this suffices.

The dependence of X and Y implies that $\text{MIC}_*(X, Y) > 0$. By Corollary 22, which states that $\sup \partial[M] = \text{MIC}_*(X, Y)$, we therefore know that there is at least one non-zero element of the boundary of $[M]$, as defined in Definition 14. Without loss of generality, suppose that this element is $[M]_{\uparrow, \ell_0} = \lim_{k \rightarrow \infty} [M]_{k, \ell_0}$. The fact that this limit is strictly positive implies that there exists some $k_0 \geq \ell_0$ and some $r > 0$ such that $[M]_{k, \ell_0} \geq r$ for all $k \geq k_0$. That is, all but finitely many of the entries in the ℓ_0 -th column of $[M]$ are at least r .

We now show that the existence of such a column suffices to prove the claim. Fix some $k > k_0$ and note that this implies that $k > \ell_0$. We argue that for all ℓ in $\{\ell_0, \dots, k\}$, the desired condition holds. Since $k > \ell_0$, the term $I^{[*]}((X, Y), k, \ell_0)$ in the definition of $[M]_{k, \ell_0}$ is a maximization over grids that have an equipartition of size k on one axis and an optimal partition of size ℓ_0 on the other. Since we allow empty rows/columns in the maximization, substituting any ℓ satisfying $\ell_0 \leq \ell \leq k$ therefore does not constrain the maximization in any way and so it cannot decrease $I^{[*]}$. In other words, for ℓ satisfying $\ell_0 \leq \ell \leq k$, we have

$$I^{[*]}((X, Y), k, \ell) \geq I^{[*]}((X, Y), k, \ell_0).$$

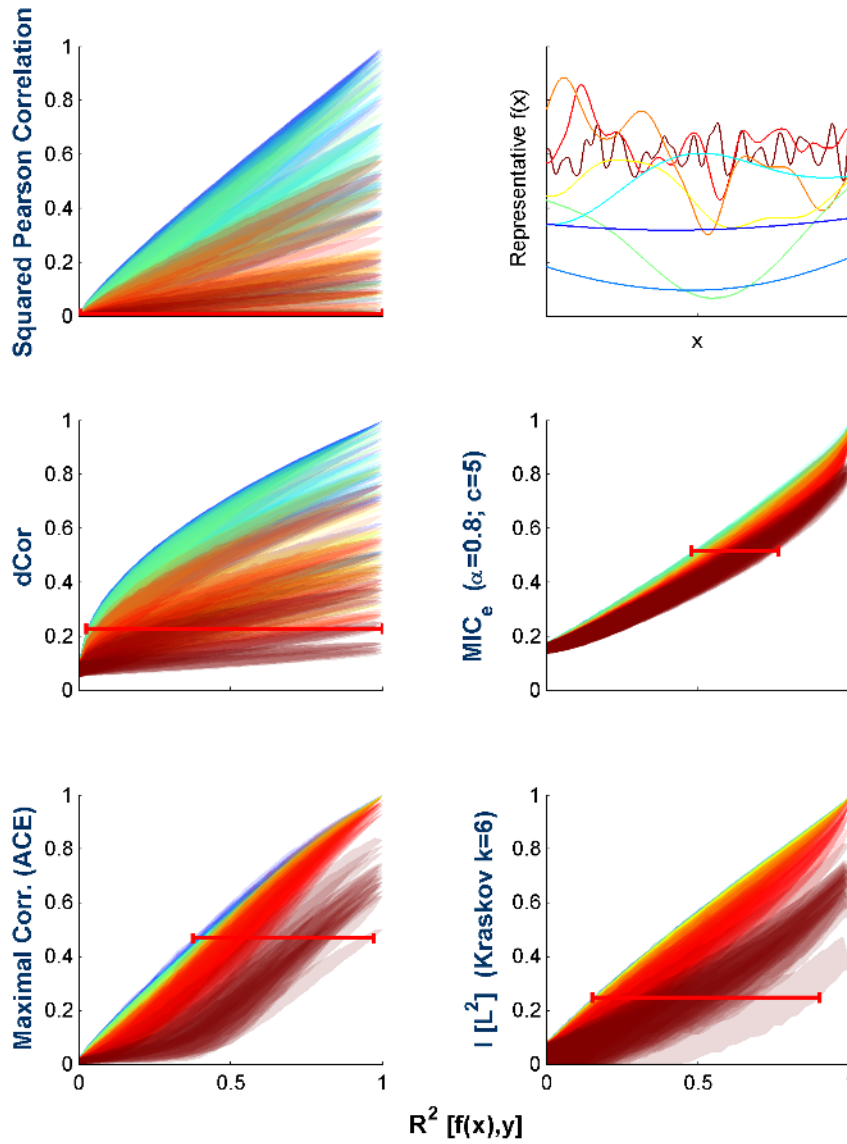


Figure J1: Equitability of methods examined on functions randomly drawn from a Gaussian process distribution, using a different noise model. This figure is identical to Figure 4, but with noise added only to the dependent variable in each relationship. Each method is assessed as in Figure 4, with a red interval indicating the widest range of R^2 values corresponding to any one value of the statistic; the narrower the red interval, the higher the equitability. Sample relationships for each Gaussian process bandwidth are shown in the top right with matching colors.

Since $k \geq \ell, \ell_0$, the normalizations in the definition of $[M]_{k,\ell}$ and $[M]_{k,\ell_0}$ are $\log \ell$ and $\log \ell_0$ respectively. Therefore, we have that

$$[M]_{k,\ell} \geq [M]_{k,\ell_0} \frac{\log \ell_0}{\log \ell} \geq \frac{r \log \ell_0}{\log \ell}$$

where the last inequality is because $k > k_0$. Setting $a = r \log \ell_0$ then gives the result. \blacksquare

K.2 Proof of Proposition 33

Proposition *Let (X, Y) be a pair of jointly distributed random variables. If X and Y are statistically independent, then $S_B(M(X, Y)) = S_B([M](X, Y)) = 0$ for all $B > 0$. If not, then $S_B(M(X, Y))$ and $S_B([M](X, Y))$ are both $\Omega(B \log \log B)$.*

Proof We give the argument for $M = M(X, Y)$ only, but the argument holds as stated for $[M](X, Y)$ as well.

The result follows from the guarantee given by the Proposition 32 above. In the case of independence, the proposition tells us that $M \equiv 0$, which immediately gives that $S_B(M) = 0$ for all $B > 0$. For the case of dependence, the proposition implies that there is some $a > 0$ and some integer $\ell_0 \geq 2$ such that, without loss of generality, $M_{k,\ell} \geq a/\log \ell$ for all $k \geq \ell \geq \ell_0$. We convert this into a lower bound on $S_B(M)$.

The key is to write the sum one column at a time, counting how many entries in each column both satisfy $k \geq \ell \geq \ell_0$ and $k\ell \leq B$. For any ℓ satisfying $\ell_0 \leq \ell \leq \sqrt{B}$, the entries $(\ell, \ell), \dots, (B/\ell, \ell)$ meet this criterion, and there are $B/\ell_0 - (\ell_0 - 1)$ of them. Moreover, since the guarantee of Proposition 32 tells us that all of these entries are at least $a/\log \ell$, we can lower-bound $S_B(M)$ as follows.

$$\begin{aligned} S_B(A) &\geq \sum_{\ell=\ell_0}^{\sqrt{B}} \frac{a}{\log \ell} \left(\frac{B}{\ell} - (\ell_0 - 1) \right) \\ &= aB \sum_{\ell=\ell_0}^{\sqrt{B}} \frac{1}{\ell \log \ell} - a \sum_{\ell=\ell_0}^{\sqrt{B}} \frac{\ell_0 - 1}{\log \ell} \\ &= a \left(B \sum_{\ell=\ell_0}^{\sqrt{B}} \frac{1}{\ell \log \ell} - O(B) \right) \\ &= \Omega(B \log \log B) \end{aligned}$$

where the second-to-last equality is because $(\ell_0 - 1)/\log \ell \leq \ell$, and the last equality is because $\sum_{i=i_0}^n 1/(i \log i)$ grows like $\log \log n$. \blacksquare

K.3 Proof of Theorem 34

Theorem *The statistics TIC_B and $TIC_{e,B}$ yield consistent right-tailed tests of independence, provided $\omega(1) < B(n) \leq O(n^{1-\varepsilon})$ for some $\varepsilon > 0$.*

Proof We give the proof for TIC only; however, the argument holds as stated for TIC_e as well.

Let (X, Y) be jointly distributed random variables, and let D_n be a sample of size n from the distribution of (X, Y) . Let $M = M(X, Y)$ be the characteristic matrix of (X, Y) and let $\widehat{M}(D_n)$ be the sample characteristic matrix. It suffices to establish the result for a deterministic monotonic function of $\text{TIC}_B(D_n)$. We therefore show convergence of $\text{TIC}_B(D_n)/B(n)$ to zero under the null hypothesis of independence and to ∞ under any alternative. Our general strategy for doing so is to translate known bounds on our error at estimating entries of M into bounds on the difference between $\text{TIC}_B(D_n)/B(n) = S_{B(n)}(\widehat{M}(D_n))/B(n)$ and $S_B(M)/B(n)$. We then obtain the result by invoking Proposition 33, which implies that $S_B(M)/B(n)$ is zero under the null hypothesis but grows without bound under the alternative.

We know from Lemma 38 (Lemma 42 for the equicharacteristic matrix) that there exists some $a > 0$ depending only on B such that

$$\left| \widehat{M}(D_n)_{k,\ell} - M_{k,\ell} \right| \leq O\left(\frac{1}{n^a}\right)$$

for all $k\ell \leq B(n)$ with probability $1 - o(1)$. This means that with probability $1 - o(1)$ we have

$$\frac{1}{B(n)} \left| \text{TIC}_B(D_n) - S_{B(n)}(M) \right| \leq O\left(\frac{\#_n}{B(n)n^a}\right)$$

where $\#_n$ is the number of pairs (k, ℓ) such that $k\ell \leq B(n)$. It can be shown by taking the integral of B/x with respect to x that $\#_n = O(B(n) \log B(n))$. Therefore, the error in the above bound is at most $O(\log B(n)/n^a) = O(1/\text{poly}(n))$ for our choice of $B(n)$.

We now use Proposition 33 to show that this bound gives the desired result. Under the null hypothesis of independence, the proposition says that $S_{B(n)}(M) = 0$ always, and so since B is a growing function the bound implies that $\text{TIC}_B(D_n)/B(n) \rightarrow 0$ in probability. Under the alternative hypothesis in which (X, Y) exhibit a dependence, the proposition implies that $S_{B(n)}(M)/B(n) > \omega(1)$. Since B is a growing function of n , this means that for any $r > 0$, the probability that $S_{B(n)}(M)/B(n) > r$ goes to 1 as n grows. In other words, $\text{TIC}_B(D_n)/B(n) \rightarrow \infty$ in probability. \blacksquare

Appendix L. Information-Theoretic Lemmas

Lemma 51 *Let Π and Ψ be random variables distributed over a discrete set of states Γ , and let (π_i) and (ψ_i) be their respective distributions. Let $P = f(\Pi)$ and $Q = f(\Psi)$ for some function f whose image is of size B . Define*

$$\varepsilon_i = \frac{\psi_i - \pi_i}{\pi_i}.$$

Then for every $0 < a < 1$ there exists some $A > 0$ such that

$$|H(Q) - H(P)| \leq (\log B) A \sum_i |\varepsilon_i|$$

when $|\varepsilon_i| \leq 1 - a$ for all i .

Proof We prove the claim with entropy measured in nats. A rescaling then gives the general result.

Let (p_i) and (q_i) be the distributions of P and Q respectively, and define

$$e_i = \frac{q_i - p_i}{p_i}$$

analogously to ε_i . Before proceeding, we observe that

$$e_i = \sum_{j \in f^{-1}(i)} \frac{\pi_j}{p_i} \varepsilon_j.$$

We now proceed with the argument. We have from [Roulston \(1999\)](#) that

$$|H(Q) - H(P)| \leq \left| \sum_i \left(e_i p_i (1 + \ln p_i) + \frac{1}{2} e_i^2 p_i + O(e_i^3) \right) \right| \quad (3)$$

$$\leq \left| \sum_i e_i p_i \right| + \left| \sum_i e_i p_i \ln p_i \right| + \frac{1}{2} \left| \sum_i e_i^2 p_i \right| + \left| \sum_i O(e_i^3) \right| \quad (4)$$

$$= \left| \sum_i e_i p_i \ln p_i \right| + \frac{1}{2} \sum_i e_i^2 p_i + \left| \sum_i O(e_i^3) \right| \quad (5)$$

where the final equality is because $\sum_i e_i p_i = \sum_i q_i - \sum_i p_i = 0$. We proceed by bounding each of the terms in Equation 5 separately.

To bound the first term, we write

$$\left| \sum_i e_i p_i \ln p_i \right| \leq - \sum_i |e_i| p_i \ln p_i.$$

We then note that $-\sum_i p_i \ln p_i \leq \ln B$, and since each of the summands has the same sign this means that $-p_i \ln p_i \leq \ln B$. We also observe that

$$|e_i| \leq \left| \sum_{j \in f^{-1}(i)} \frac{\pi_j}{p_i} \varepsilon_j \right| \leq \sum_j \frac{\pi_j}{p_i} |\varepsilon_j| \leq \sum_j |\varepsilon_j|$$

since $\pi_j/p_i \leq 1$. Together, these two facts give

$$\begin{aligned} - \sum_i |e_i| p_i \ln p_i &\leq (\ln B) \sum_i |e_i| \\ &\leq (\ln B) \sum_i |\varepsilon_i| \end{aligned}$$

The second inequality is because each e_i is a weighted average of a set of ε_i and each ε_i enters into the expression of exactly one e_i .

To bound the second term, we use the fact that $p_i \leq 1$ for all i , and so

$$\sum_i e_i^2 p_i \leq \sum_i e_i^2.$$

We then write

$$\begin{aligned}
\sum_i e_i^2 &= \sum_i \left(\sum_{j \in f^{-1}(i)} \frac{\pi_j}{p_i} \varepsilon_j \right)^2 \\
&\leq \sum_i \sum_{j \in f^{-1}(i)} \frac{\pi_j}{p_i} \varepsilon_j^2 \\
&\leq \sum_j \varepsilon_j^2 \\
&= \sum_j O(|\varepsilon_j|)
\end{aligned}$$

where the second line is a consequence of the convexity of $f(x) = x^2$ and the third line is because the sets $f^{-1}(i)$ partition Γ .

To bound the third term, we write

$$\left| \sum_i O(e_i^3) \right| \leq \sum_i O(|e_i|^3)$$

and then proceed as we did with the second term, using the fact that $f(x) = x^3$ is convex for $x \geq 0$. This gives

$$\sum_i O(|e_i|^3) \leq \sum_i O(|\varepsilon_i|^3) = \sum_i O(|\varepsilon_i|)$$

completing the proof. ■

Lemma 52 *Let $\{w_i\} \subset [0, 1]$ be a set of size n with $\sum_i w_i \leq 1$, and let $\{u_i\}$ be a set of n non-negative numbers satisfying $\sum_i u_i = a$ and $u_i \leq w_i$. Then*

$$\sum_{i=1}^n w_i H_b \left(\frac{u_i}{w_i} \right) \leq H_b(a)$$

where H_b is the binary entropy function.

Proof Consider the random variable X taking values in $\{0, \dots, n\}$ that equals zero with probability $1 - \sum_i w_i$ and equals i with probability w_i for $0 < i \leq n$. Define the random variable Y taking values in $\{0, 1\}$ by

$$\mathbf{P}(Y = 0 | X = i) = \begin{cases} 0 & i = 0 \\ u_i/w_i & 0 < i \leq n \end{cases} .$$

The function we wish to bound equals $H(Y|X) \leq H(Y)$. We therefore observe that

$$\sum_{i=1}^n w_i H_b \left(\frac{u_i}{w_i} \right) \leq H(Y).$$

The result follows from the observation that

$$\mathbf{P}(Y = 0) = \sum_i \mathbf{P}(X = i) \frac{u_i}{w_i} = \sum_i u_i \leq a.$$

■

Lemma 53 *Let X be a random variable distributed over k states, with $\mathbf{P}(X = x) = p_x$. Let $\alpha_x \geq 0$ be such that $\sum \alpha_x = \delta$, and define the random variable X' by $\mathbf{P}(X' = x) = (p_x + \alpha_x)/(1 + \delta)$. We have*

$$|H(X') - H(X)| \leq H_b(\delta) + \delta \log k$$

where H_b is the binary entropy function.

Proof Define a new random variable Z by

$$\mathbf{P}(Z = 0|X' = x) = \frac{p_x}{p_x + \alpha_x}, \quad \mathbf{P}(Z = 1|X' = x) = \frac{\alpha_x}{p_x + \alpha_x}.$$

We will use the fact that $H(X'|Z = 0) = H(X)$ to obtain the required bound.

To upper bound $H(X') - H(X)$, we write

$$\begin{aligned} H(X') - H(X) &\leq H(X', Z) - H(X) \\ &= H(Z) + \mathbf{P}(Z = 0)H(X'|Z = 0) + \mathbf{P}(Z = 1)H(X'|Z = 1) - H(X) \\ &\leq H_b(\delta) + (1 - \delta)H(X) + \delta H(X'|Z = 1) - H(X) \\ &= H_b(\delta) - \delta H(X) + \delta \log k \\ &\leq H_b(\delta) + \delta \log k \end{aligned}$$

where in the fourth line we have used that $H(X'|Z = 1) \leq \log k$.

To upper bound $H(X) - H(X')$, we write

$$\begin{aligned} H(X') + H(Z) &\geq H(X', Z) \\ &\geq \mathbf{P}(Z = 0)H(X'|Z = 0) \\ &= (1 - \delta)H(X) \end{aligned}$$

which yields

$$H(X') \geq (1 - \delta)H(X) - H_b(\delta)$$

since $H(Z) = H_b(\delta)$. Thus, we have

$$H(X) - H(X') \leq \delta H(X) + H_b(\delta) \leq \delta \log k + H_b(\delta).$$

■

Lemma 54 *Let X be a random variable distributed over k states, with $\mathbf{P}(X = x) = p_x$. Let $\alpha_x \leq 0$ be such that $\sum |\alpha_x| = \delta$, and define the random variable X' by $\mathbf{P}(X' = x) = (p_x + \alpha_x)/(1 - \delta)$. We have*

$$|H(X') - H(X)| \leq H_b\left(\frac{\delta}{1 - \delta}\right) + \frac{\delta}{1 - \delta} \log k$$

where H_b is the binary entropy function. In particular, when $\delta \leq 1/3$ we have

$$|H(X') - H(X)| \leq H_b(2\delta) + 2\delta \log k.$$

Proof We observe that we can get from X' to X by adding $\delta/(1 - \delta)$ probability mass and rescaling. The previous lemma then gives the result. \blacksquare

Lemma 55 *Let X be a random variable distributed over k states, with $\mathbf{P}(X = x) = p_x$. Let α_x be such that $\sum |\alpha_x| = \delta$, and define the random variable X' by $\mathbf{P}(X' = x) = (p_x + \alpha_x)/(1 - \sum \alpha_x)$. That is, X' is the result of changing the probability of state x by α_x and then re-normalizing to obtain a valid distribution. If $\delta \leq 1/4$, we have*

$$|H(X') - H(X)| \leq 2H_b(2\delta) + 3\delta \log k$$

where H_b is the binary entropy function.

Proof Let δ_+ be the total magnitude of all the positive α_x , and let δ_- be the total magnitude of all the negative α_x . We first add all the mass we're going to add, and apply the first of the previous two lemmas. Then we remove all the mass we are going to remove, and apply the second of the two previous lemmas. This yields a bound of

$$\begin{aligned} & H_b(\delta_+) + \delta_+ \log k + H_b\left(2\frac{\delta_-}{1 + \delta_+}\right) + 2\frac{\delta_-}{1 + \delta_+} \log k \\ & \leq H_b(\delta_+) + \delta_+ \log k + H_b(2\delta_-) + 2\delta_- \log k \\ & \leq H_b(2\delta) + \delta \log k + H_b(2\delta) + 2\delta \log k \\ & \leq 2H_b(2\delta) + 3\delta \log k \end{aligned}$$

where the first inequality is because $1 + \delta_+ \leq 1 + \delta < 2$ and $2\delta_- \leq 2\delta \leq 1/2$, and the second inequality is because $\delta_+ \leq \delta < 2\delta \leq 1/2$. \blacksquare

References

- Leo Breiman and Jerome H. Friedman. Estimating optimal transformations for multiple regression and correlation. *Journal of the American statistical Association*, 80(391):580–598, 1985.
- William S. Cleveland and Susan J. Devlin. Locally weighted regression: an approach to regression analysis by local fitting. *Journal of the American Statistical Association*, 83(403):596–610, 1988.
- Thomas Cover and Joy Thomas. *Elements of Information Theory*. New York: John Wiley & Sons, Inc, 2006.
- Imre Csiszár. Axiomatic characterizations of information measures. *Entropy*, 10(3):261–273, 2008.
- Imre Csiszár and Paul C. Shields. Information theory and statistics: A tutorial. *Communications and Information Theory*, 1(4):417–528, 2004.
- A. Adam Ding and Yi Li. Copula correlation: An equitable dependence measure and extension of pearson’s correlation. *arXiv preprint arXiv:1312.7214*, 2013.
- Valur Emilsson, Gudmar Thorleifsson, Bin Zhang, Amy S Leonardson, Florian Zink, Jun Zhu, Sonia Carlson, Agnar Helgason, G Bragi Walters, Steinunn Gunnarsdottir, et al. Genetics of gene expression and its effect on disease. *Nature*, 452(7186):423–428, 2008.
- 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.
- Arthur Gretton, Kenji Fukumizu, Choon H Teo, Le Song, Bernhard Schölkopf, and Alex J Smola. A kernel statistical test of independence. In *Advances in neural information processing systems*, pages 585–592, 2007.
- Arthur Gretton, Karsten M. Borgwardt, Malte J. Rasch, Bernhard Schölkopf, and Alexander Smola. A kernel two-sample test. *The Journal of Machine Learning Research*, 13(1):723–773, 2012.
- Ruth Heller, Yair Heller, and Malka Gorfine. A consistent multivariate test of association based on ranks of distances. *Biometrika*, 100(2):503–510, 2013.
- Ruth Heller, Yair Heller, Shachar Kaufman, Barak Brill, and Malka Gorfine. Consistent distribution-free k -sample and independence tests for univariate random variables. *Journal of Machine Learning Research*, 17(29):1–54, 2016.
- Wassily Hoeffding. A non-parametric test of independence. *The Annals of Mathematical Statistics*, pages 546–557, 1948.
- Bo Jiang, Chao Ye, and Jun S Liu. Nonparametric k -sample tests via dynamic slicing. *Journal of the American Statistical Association*, 110(510):642–653, 2015.

- Justin B. Kinney and Gurinder S. Atwal. Equitability, mutual information, and the maximal information coefficient. *Proceedings of the National Academy of Sciences*, 2014. doi: 10.1073/pnas.1309933111.
- Alexander Kraskov, Harald Stogbauer, and Peter Grassberger. Estimating mutual information. *Physical Review E*, 69, 2004.
- Edward H. Linfoot. An informational measure of correlation. *Information and Control*, 1(1):85–89, 1957.
- David Lopez-Paz, Philipp Hennig, and Bernhard Schölkopf. The randomized dependence coefficient. In *Advances in Neural Information Processing Systems*, pages 1–9, 2013.
- Michael Mitzenmacher and Eli Upfal. *Probability and computing: Randomized algorithms and probabilistic analysis*. Cambridge University Press, 2005.
- Ben Murrell, Daniel Murrell, and Hugh Murrell. R2-equitability is satisfiable. *Proceedings of the National Academy of Sciences*, 2014. doi: 10.1073/pnas.1403623111. URL <http://www.pnas.org/content/early/2014/04/29/1403623111.short>.
- Liam Paninski. Estimation of entropy and mutual information. *Neural computation*, 15(6):1191–1253, 2003.
- Alfred Rényi. On measures of dependence. *Acta mathematica hungarica*, 10(3):441–451, 1959.
- David N. Reshef, Yakir A. Reshef, Hilary K. Finucane, Sharon R. Grossman, Gilean McVean, Peter J. Turnbaugh, Eric S. Lander, Michael Mitzenmacher, and Pardis C. Sabeti. Detecting novel associations in large data sets. *Science*, 334(6062):1518–1524, 2011.
- David N. Reshef, Yakir A. Reshef, Michael Mitzenmacher, and Pardis C. Sabeti. Cleaning up the record on the maximal information coefficient and equitability. *Proceedings of the National Academy of Sciences*, 2014. doi: 10.1073/pnas.1408920111. URL <http://www.pnas.org/content/early/2014/08/07/1408920111.short>.
- David N. Reshef, Yakir A. Reshef, Pardis C. Sabeti, and Michael Mitzenmacher. An empirical study of leading measures of dependence. *arXiv preprint arXiv:1505.02214*, 2015a.
- Yakir A Reshef, David N Reshef, Pardis C Sabeti, and Michael Mitzenmacher. Equitability, interval estimation, and statistical power. *arXiv preprint arXiv:1505.02212*, 2015b.
- Mark S. Roulston. Estimating the errors on measured entropy and mutual information. *Physica D: Nonlinear Phenomena*, 125(3):285–294, 1999.
- Noah Simon and Robert Tibshirani. Comment on “Detecting novel associations in large data sets”. *Unpublished (available at <http://www-stat.stanford.edu/~tibs/reshef/comment.pdf> on 11 Nov. 2012)*, 2012.
- Terry Speed. A correlation for the 21st century. *Science*, 334(6062):1502–1503, 2011.
- Josef Stoer and Roland Bulirsch. *Introduction to Numerical Analysis*. Springer-Verlag, 1980.

Charles J. Stone. Consistent nonparametric regression. *The annals of statistics*, pages 595–620, 1977.

John D. Storey and Robert Tibshirani. Statistical significance for genomewide studies. *Proceedings of the National Academy of Sciences*, 100(16):9440–9445, 2003.

Gábor J. Székely and Maria L. Rizzo. Brownian distance covariance. *The Annals of Applied Statistics*, 3(4):1236–1265, 2009.

Gábor J. Székely, Maria L. Rizzo, Nail K. Bakirov, et al. Measuring and testing dependence by correlation of distances. *The Annals of Statistics*, 35(6):2769–2794, 2007.