

Discovering and Deciphering Relationships Across Disparate Data Modalities

Joshua T. Vogelstein^{1,2}, Eric W. Bridgeford¹, Qing Wang¹, Carey E. Priebe¹, Mauro Maggioni¹, and Cencheng Shen^{1,3}

¹Johns Hopkins University

²Child Mind Institute

³University of Delaware

Understanding the relationships between different properties of data, such as whether a genome or connectome has information about disease status, is increasingly important. While existing approaches can test whether two properties are related, they may require unfeasibly large sample sizes and often are not interpretable. Our approach, “Multiscale Graph Correlation” (MGC), is a dependence test that juxtaposes disparate data science techniques, including k-nearest neighbors, kernel methods, and multiscale analysis. Other methods may require double or triple the number of samples to achieve the same statistical power as MGC in a benchmark suite including high-dimensional and nonlinear relationships, with dimensionality ranging from 1 to 1000. Moreover, MGC uniquely characterizes the latent geometry underlying the relationship, while maintaining computational efficiency. In real data, including brain imaging and cancer genetics, MGC detects the presence of a dependency and provides guidance for the next experiments to conduct.

Introduction

Identifying the existence of a relationship between a pair of properties or modalities is the critical initial step in data science investigations. Only if there is a statistically significant relationship does it make sense to try to decipher the nature of the relationship. Discovering and deciphering relationships is fundamental, for example, in high-throughput screening [1], precision medicine [2], machine learning [3], and causal analyses [4]. One of the first approaches for determining whether two properties are related to—or statistically dependent on—each other is Pearson’s Product-Moment Correlation (published in 1895 [5]). This seminal paper prompted the development of entirely new ways of thinking about and quantifying relationships (see [6, 7] for recent reviews and discussion). Modern datasets, however, present challenges for dependence-testing that were not addressed in Pearson’s era. First, we now desire methods that can correctly detect any kind of dependence between all kinds of data, including high-dimensional data (such as ‘omics), structured data (such as images or networks), with nonlinear relationships (such as oscillators), even with very small sample sizes as is common in modern biomedical science. Second, we desire methods that are interpretable by providing insight into how or why they discovered the presence of a statistically significant relationship. Such insight can be a crucial component of designing the next computational or physical experiment.

While many statistical and machine learning approaches have been developed over the last 120 years to combat aspects of the first issue—detecting dependencies—no approach satisfactorily addressed the challenges across all data types, relationships, and dimensionalities. Hoeffding and Renyi proposed non-parametric tests to address nonlinear but univariate relationships [8, 9]. In the 1970s and 1980s, nearest neighbor style approaches were popularized [10, 11], but they were sensitive to algorithm parameters resulting in poor empirical performance. “Energy statistics”, and in particular the distance correlation test (DCORR), was recently shown to be able to detect any dependency with sufficient

observations, at arbitrary dimensions, and structured data under a proper distance metric [12–15]. Another set of methods, referred to as “kernel mean embedding” approaches, including the Hilbert Schmidt Independence Criterion (HSIC) [16, 17], have the same theoretical guarantees, which is shown to be a kernel version of the energy statistics [18, 19]. The energy statistics can perform very well with a relatively small sample size on high-dimensional linear data, whereas the kernel methods and another test (Heller, Heller, and Gorfine’s test, H_{HG}) [20] perform well on low-dimensional nonlinear data. But no test performs particularly well on high-dimensional nonlinear data with typical sample sizes, which characterizes a large fraction of real data challenges in the current big data era.

Moreover, to our knowledge, existing dependency tests do not attempt to further characterize the dependency structure. On the other hand, much effort has been devoted to characterizing “point cloud data”, that is, summarizing certain global properties in unsupervised settings (for example, having genomics data, but no disease data). Classic examples of such approaches include Fourier [21] and wavelet analysis [22]. More recently, topological and geometric data analysis compute properties of graphs, or even higher order simplices [23]. Such methods build multiscale characterization of the samples, much like recent developments in harmonic analysis [24, 25]. However, these tools typically lack statistical guarantees under noisy observations, and are often computationally burdensome.

We surmised that both (i) empirical performance in all dependency structures, in particular high-dimensional, nonlinear, low-sample size settings, and (ii) providing insight into the discovery process, can be addressed via extending existing dependence tests to be *adaptive* to the data [26]. Existing tests rely on a fixed *a priori* selection of an algorithmic parameter, such as the kernel bandwidth [27], intrinsic dimension [25], and/or local scale [10, 11]. Indeed, the Achilles Heel of manifold learning has been the requirement to manually choose these parameters [28]. Post-hoc cross-validation is often used to make these methods effectively adaptive, but doing so adds an undesirable computational burden, and may weaken or destroy any statistical guarantees. There is therefore a need for statistically valid and computationally efficient adaptive methods.

To illustrate the importance of adapting to different kinds of relationships, consider a simple illustrative example: investigate the relationship between cloud density and grass wetness. If this relationship were approximately linear, the data might look like those in Figure 1A (top). On the other hand, if the relationship were nonlinear—such as a spiral—it might look like those in Figure 1A (bottom). Although the relationship between clouds and grass is unlikely to be spiral, spiral relationships are prevalent in nature and mathematics (for example, shells, hurricanes, and galaxies), and are canonical in evaluations of manifold learning techniques [29], thereby motivating its use here.

Under the linear relationship (top panels), when a pair of observations are close to each other in cloud density, they also tend to be close to each other in grass wetness (for example, observations 1 and 2 highlighted in black in Figure 1A, and distances between them in Figure 1B). Similarly, when a pair of observations are far from each other in cloud density, they also tend to be far from each other in grass wetness (see for example, distances between observations 2 and 3). On the other hand, consider the nonlinear (spiral) relationship (bottom panels). Here, when a pair of observations are close to each other in cloud density, they also tend to be close to each other in grass wetness (see points 1 and 2 again). However, the same is not true for large distances (see points 2 and 3). Thus, in the linear relationship, the distance between every pair of points is informative with respect to the relationship, while under the nonlinear relationship, only a subset of the distances are.

For this reason, we juxtapose nearest neighbor mechanism with distance methods. Specifically, for each point, we find its k -nearest neighbors for one property (e.g., cloud density), and its l -nearest neighbors for the other property (e.g., grass wetness); we call the pair (k, l) the “scale”. *A priori*, however, we do not know which scales will be most informative. We compute all distance pairs, then ef-

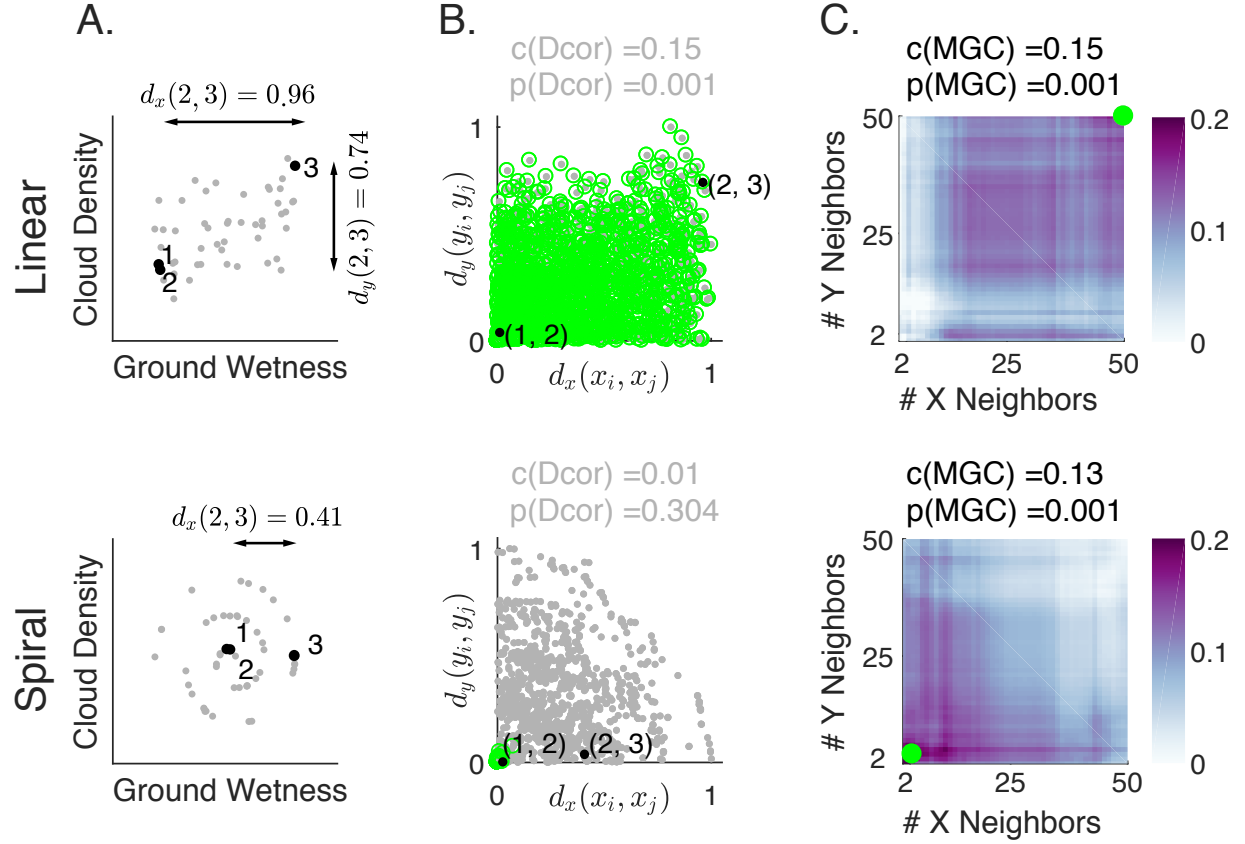


Figure 1: Illustration of Multiscale Graph Correlation (Mgc) on simulated cloud density (x_i) and grass wetness (y_i). We present two different relationships: linear (top) and nonlinear spiral (bottom; see Materials and Methods for simulation details). **(A)** Scatterplots of the raw data using 50 pairs of samples for each scenario. Samples 1, 2, and 3 (black) are highlighted; arrows show x distances between these pairs of points while their y distances are almost 0. **(B)** Scatterplots of all pairs of distances comparing x and y distances. Distances are linearly correlated in the linear relationship, whereas they are not in the spiral relationship. Dcorr uses all distances (gray dots) to compute its test statistic and p-value, whereas Mgc chooses the local scale and then uses only the local distances (green dots). **(C)** Heatmaps characterizing the strength of the generalized correlation at all possible scales (ranging from 2 to n for both x and y). For the linear relationship, the global scale is optimal, which is the scale that Mgc selects and results in a p-value identical to Dcorr. For the nonlinear relationship, the optimal scale is local in both x and y , so Mgc achieves a far larger test statistic, and a correspondingly smaller and significant p-value. Thus, Mgc uniquely detects dependence and characterizes the geometry in both relationships.

ficiently compute the distance correlations for all scales. The local correlations (Figure 1C, described in detail below) illustrate which scales are relatively informative about the relationship. The key, therefore, to successfully discover and decipher relationships between disparate data modalities is to adaptively determine which scales are the most informative, and the geometric implication for the most informative scales. Doing so not only provides an estimate of whether the modalities are related, but also provides insight into how the determination was made. This is especially important in high-dimensional data, where simple visualizations do not reveal relationships to the unaided human eye.

Our method, “Multiscale Graph Correlation” (MGC, pronounced “magic”), generalized and extends previously proposed pairwise comparison-based approaches by adaptively estimating the informative scales for any relationship—linear or nonlinear, low-dimensional or high-dimensional, unstructured or structured—in a computationally efficient and statistically valid and consistent fashion. This adaptive nature of MGC effectively guarantees an improved statistical performance. Moreover, the dependency strength across all scales is informative about the structure of a statistical relationship, therefore providing further guidance for subsequent experimental or analytical steps. MGC is thus a hypothesis-testing and insight-providing approach that builds on recent developments in manifold and kernel learning, with complementary developments in nearest-neighbor search, and multiscale analyses.

The Multiscale Graph Correlation Procedure

MGC is a multi-step procedure to discover and decipher dependencies across disparate data modalities or properties. Given n samples of two different properties, proceed as follows (see Materials and Methods and Shen et al. [30] for details):

1. Compute two distance matrices, one consisting of distances between all pairs of one property (e.g., cloud densities, entire genomes or connectomes) and the other consisting of distances between all pairs of the other property (e.g., grass wetnesses or disease status). Then center each distance matrix (by subtracting its overall mean, the column-wise mean from each column, and the row-wise mean from each row), and denote the resulting n -by- n matrices A and B .
2. For all possible values of k and l from 1 to n :
 - (a) Compute the k -nearest neighbor graphs for one property, and the l -nearest neighbor graphs for the other property. Let G_k and H_l be the adjacency matrices for the nearest neighbor graphs, so that $G_k(i, j) = 1$ indicates that $A(i, j)$ is within the k smallest values of the i^{th} row of A , and similarly for H_l .
 - (b) Estimate the local correlations—the correlation between distances restricted to only the (k, l) neighbors—by summing the products of the above matrices,
$$c^{kl} = \sum_{ij} A(i, j)G_k(i, j)B(i, j)H_l(i, j).$$
 - (c) Normalize c^{kl} such that the result is always between -1 and $+1$ by dividing by
$$\sqrt{\sum_{ij} A^2(i, j)G_k(i, j) \times \sum_{ij} B^2(i, j)H_l(i, j)}.$$
3. Estimate the optimal local correlation c^* by finding the smoothed maximum of all local correlations $\{c^{kl}\}$. Smoothing mitigates biases and provides MGC with theoretical guarantee and better finite-sample performance.
4. Determine whether the relationship is significantly dependent—that is, whether c^* is more extreme than expected under the null—via a permutation test. The permutation procedure repeats steps 1-4 on each permutation, thereby eliminating the multiple hypothesis testing problem by

only computing one overall p-value, rather than one p-value per scale, ensuring that it is a valid test (meaning that the false positive rate is properly controlled at the specified type I error rate).

Computing all local correlations, the test statistic, and the p-value requires $O(n^2 \log n)$ time, which is about the same running time complexity as other methods [30].

Results

MGC Typically Requires Substantially Fewer Samples to Achieve the Same Power Across All Dependencies and Dimensions

When, and to what extent, does MGC outperform other approaches, and when does it not? To address this question, we formally pose the following hypothesis test (see Materials and Methods for details):

$$H_0: X \text{ and } Y \text{ are independent}$$

$$H_A: X \text{ and } Y \text{ are not independent.}$$

The standard criterion for evaluating statistical tests is the testing power, which equals the probability that a test correctly rejects the null hypothesis at a given type 1 error level, i.e., $\text{power} = \text{Prob}(H_0 \text{ is rejected} \mid H_0 \text{ is false})$. The higher the testing power, the better the test procedure. A consistent test has power converging to 1 under dependence, and a valid test controls the type 1 error level under independence. In a complementary manuscript [30], we established the theoretical properties of MGC, proving its validity and universal consistency for dependence testing against all distributions of finite second moments.

Here, we address the empirical performance of MGC as compared with multiple popular tests: (i) Dcorr, a popular approach from the statistics community [12, 13], (ii) Mcorr, a modified version of Dcorr designed to be unbiased for sample data [14], (iii) Hhg, a distance-based test that is very powerful for detecting low-dimensional nonlinear relationships [20], (iv) Hsic, a kernel dependency measure [16] formulated in the same way as Dcorr except operating on kernels, (v) Mantel, which is historically widely used in biology and ecology [31], (vi) RV coefficient [5, 7], which is a multivariate generalization of PEARSON's product moment correlation whose test statistic is the sum of the trace-norm of the cross-covariance matrix, and (vii) the CCA method, which is the largest (in magnitude) singular value of the cross-covariance matrix, and can be viewed as a different generalization of PEARSON in high-dimensions that is more appropriate for sparse settings [32–34]. Note that while we focus on high-dimensional settings, Appendix shows further results in one-dimensional settings, also comparing to a number of tests that are limited to one dimension, including: (viii) PEARSON's product moment correlation, (ix) SPEARMAN's rank correlation [35], (x) KENDALL's tau correlation [36], and (xi) MIC [37]. Under the regularity condition that the data distribution has finite second moment, the first four tests are universally consistent, whereas the other tests are not.

We generate an extensive benchmark suite of 20 relationships, including different polynomial (linear, quadratic, cubic), trigonometric (sinusoidal, circular, ellipsoidal, spiral), geometric (square, diamond, W-shape), and other functions. This suite includes and extends the simulated settings from previous dependence testing work [12, 14, 20, 38, 39]. For many of them, we introduce high-dimensional variants, to more extensively evaluate the methods; function details are in Materials and Methods. The visualization of one-dimensional noise-free (black) and noisy (gray) samples is shown in Figure 2- Figure Supplement 1. For each relationship, we compute the power of each method relative to MGC for ~ 20 different dimensionalities, ranging from 1 up to 10, 20, 40, 100, or 1000. The high-dimensional relationships are more challenging because (1) they cannot be easily visualized, and (2) each dimension is designed to have less and less signal, so there are many noisy dimensions. Figure 2 shows

that Mgc achieves the highest (or close to the highest) power given 100 samples for each relationship and dimensionality. Figure 2-Figure Supplement 2 shows the same advantage in one-dimension with increasing sample size.

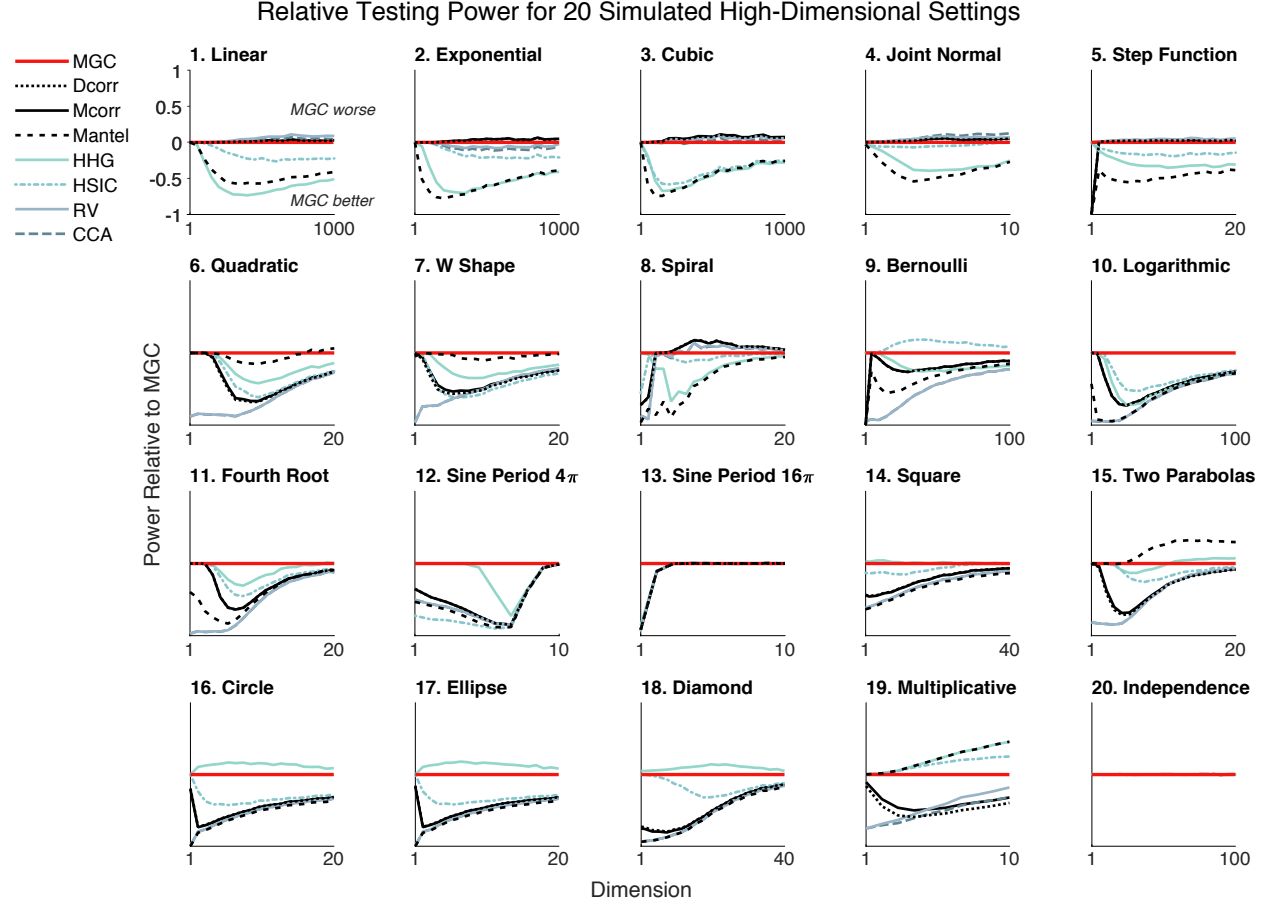


Figure 2: An extensive benchmark suite of 20 different relationships spanning polynomial, trigonometric, geometric, and other relationships demonstrates that Mgc empirically nearly dominates eight other methods across dependencies and dimensionalities ranging from 1 to 1000 (see Material and Methods and Figure 2-Figure Supplement 1 for details). Each panel shows the testing power of other methods relative to the power of Mgc (e.g., power of Mgc minus the power of Mgc) at significance level $\alpha = 0.05$ versus dimensionality for $n = 100$. Any line below zero at any point indicates that that method's power is less than Mgc's power for the specified setting and dimensionality. Mgc achieves empirically better (or similar) power than all other methods in almost all relationships and all dimensions. For the the independent relationship (#20), all methods yield power 0.05 as they should. Note that Mgc is always plotted "on top" of the other methods, therefore, some lines are obscured.

Moreover, for each relationship and each method we compute the required sample size to achieve power 85% at error level 0.05, and summarize the median size for monotone relationships (type 1-5) and non-monotone relationships (type 6-19) in Table 1. Other methods typically require double or triple the number of samples as Mgc to achieve the same power. More specifically, traditional correlation methods (PEARSON, RV, CCA, SPEARMAN, KENDALL) always perform the best in monotonic simulations, distance-based methods including Mgc, Dcorr, Mgc, HHG and HSIC are slightly worse, while MIC and MANTEL are the worst. Mgc's performance is equal to linear methods on monotonic relationships. For non-monotonic relationships, traditional correlations fail to detect the existence of dependencies, Dcorr, Mgc, and MIC, do reasonably well, but HHG and Mgc require the fewest sam-

ples. In the high-dimensional non-monotonic relationships that motivated this work, and are common in biomedicine, Mgc significantly outperforms other methods. The second best test that is universally consistent (HHG) requires nearly double as many samples as Mgc, demonstrating that Mgc could half the time and cost of experiments designed to discover relationships at a given effect size.

Mgc extends previously proposed global methods, such as MANTEL and DCORR. The above experiments extended MCORR, because MCORR is universally consistent and an unbiased version of DCORR [14]. Figure 2-Figure Supplement 3 directly compares multiscale generalizations of MANTEL and MCORR as dimension increases, demonstrating that empirically, Mgc nearly dominates its global variant for essentially all dimensions and simulation settings considered here. Figure 2-Figure Supplement 4 shows a similar result for one-dimensional settings while varying sample size. Thus, not only does Mgc empirically nearly dominate existing tests, it is a framework that one can apply to future tests to further improve their performance.

Table 1: The median sample size for each method to achieve power 85% at type 1 error level 0.05, grouped into monotone (type 1-5) and non-monotone relationships (type 6-19) for both one- and ten-dimensional settings, normalized by the number of samples required by Mgc. In other words, a 2.0 indicates that the method requires double the sample size to achieve 85% power relative to Mgc. PEARSON, RV, and CCA all achieve the same performance, as do SPEARMAN and KENDALL. Mgc requires the fewest number of samples in all settings, and for high-dimensional non-monotonic relationships, all other methods require about double or triple the number of samples Mgc requires.

Dimensionality Dependency Type	One-Dimensional			Ten-Dimensional		
	Monotone	Non-Mono	Average	Monotone	Non-Mono	Average
Mgc	1	1	1	1	1	1
DCORR	1	2.6	2.2	1	3.2	2.6
MCORR	1	2.8	2.4	1	3.1	2.6
HHG	1.4	1	1.1	1.7	1.9	1.8
HSIC	1.4	1.1	1.2	1.7	2.4	2.2
MANTEL	1.4	1.8	1.7	3	1.6	1.9
PEARSON / RV / CCA	1	>10	>10	0.8	>10	>10
SPEARMAN / KENDALL	1	>10	>10	n/a	n/a	n/a
MIC	2.4	2	2.1	n/a	n/a	n/a

Mgc Deciphers Latent Dependence Structure

Beyond simply testing the existence of a relationship, the next goal is often to decipher the nature or structure of the relationship, thereby providing insight and guiding future experiments. A single scalar quantity (such as effect size) is inadequate given the vastness and complexities of possible relationships. Existing methods would require a secondary procedure to characterize the relationship, which introduces complicated “post selection” statistical quandaries that remain mostly unresolved [40]. Instead, Mgc provides a simple, intuitive, and nonparametric (and therefore infinitely flexible) “map” of how it discovered the relationship. As described below, this map not only provides interpretability for how Mgc detected a dependence, it also partially characterizes the geometry of the investigated relationship.

The Mgc-Map shows local correlation as a function of the scales of the two properties. More concretely, it is the matrix of c^{kl} s, as defined above. Thus, the Mgc-Map is an n-by-n matrix which encodes the strength of dependence for each possible scale. Figure 3 provides the Mgc-Map for all 20 different one-dimensional relationships; the optimal scale to achieve \hat{t}_* is marked with a green dot. For the

MGC Images Characterize the Geometry of Dependence

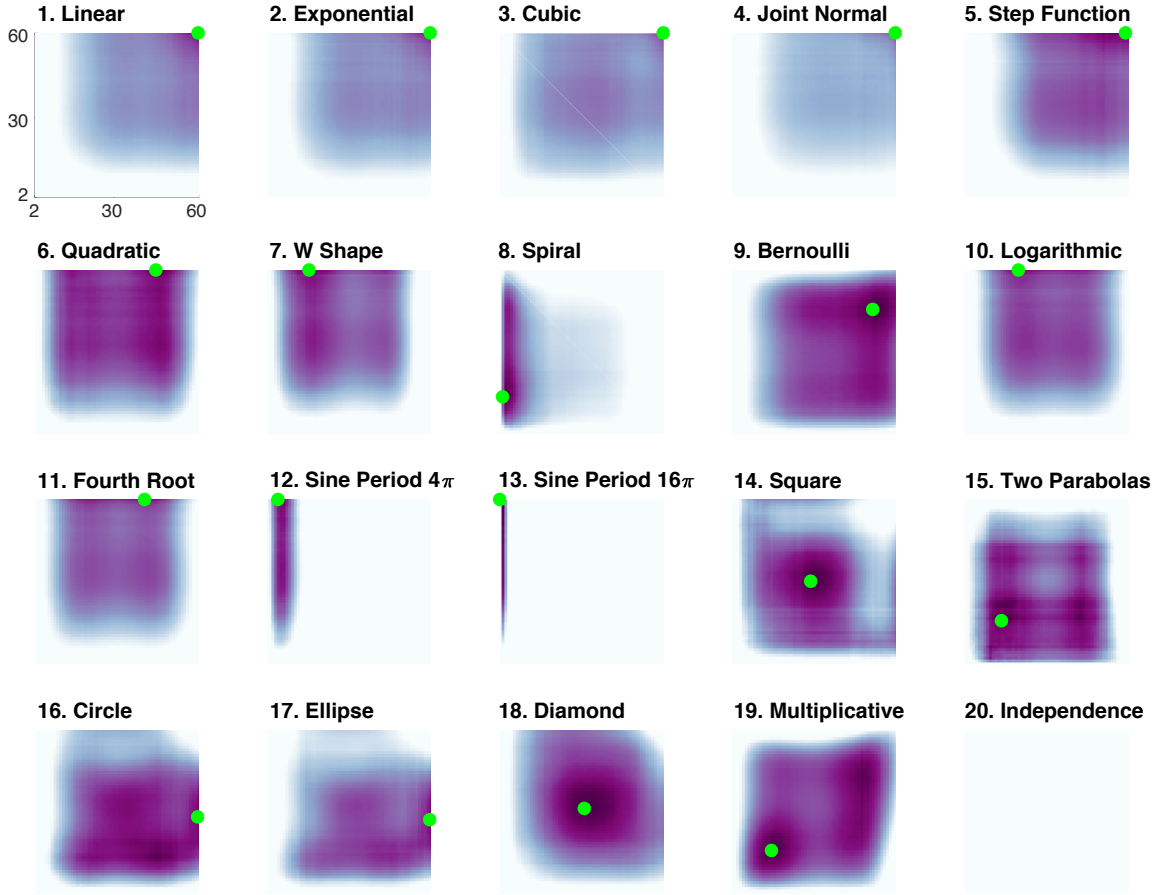


Figure 3: The Mgc-Map characterizes the geometry of the dependence function. For each of the 20 panels, the abscissa and ordinate denote the number of neighbors for X and Y , respectively, and the color denotes the magnitude of each local correlation. For each simulation, the sample size is 60, and both X and Y are one-dimensional. Each dependency has a different Mgc-Map characterizing the geometry of dependence, and the optimal scale is shown in green. In linear or close-to-linear relationships (first row), the optimal scale is global, i.e., the green dot is in the top right corner. Otherwise the optimal scale is non-global, which holds for the remaining dependencies. Moreover, similar dependencies often share similar Mgc-Maps and similar optimal scales, such as (10) logarithmic and (11) fourth root, the trigonometric functions in (12) and (13), (16) circle and (17) ellipse, and (14) square and (18) diamond. The Mgc-Maps for high-dimensional simulations are provided in Figure 3-Figure Supplement 1.

monotonic dependencies (1-5), the optimal scale is always the largest scale, i.e., the global one. For all non-monotonic dependencies (6-19), Mgc chooses smaller scales. Thus, a global optimal scale implies a close-to-linear dependency, otherwise the dependency is strongly nonlinear. In fact, this empirical observation led to the following theorem (which is proved in Materials and Methods) :

Theorem 1. *When (X, Y) are linearly related (meaning that Y can be constructed from X by rotation, scaling, translation, and/or reflection), the optimal scale of Mgc equals the global scale. Conversely, a local optimal scale implies a nonlinear relationship.*

Thus, the Mgc-Map explains how Mgc discovers relationships, specifically, which scale has the most

informative pairwise comparisons, and how that relates to the geometry of the relationship. Note that Mgc provides the geometric characterization “for free”, meaning that no separate procedure is required; therefore, Mgc provides both a valid test and information about the geometric relationship. Moreover, similar dependencies have similar Mgc-Maps and often similar optimal scales. For example, logarithmic (10) and fourth root (11), though very different functions analytically, are geometrically similar, and yield very similar Mgc-Maps. Similarly, (12) and (13) are trigonometric functions, and they share a narrow range of significant local scales. Both circle (16) and ellipse (17), as well as square (14) and diamond (18), are closely related geometrically and also have similar Mgc-Maps. This indicates that the Mgc-Map partially characterizes the geometry of these relationships, differentiating different dependence structures and assisting subsequent analysis steps. Moreover, in [30] we proved that the sample Mgc-Map (which Mgc estimates) converges to the true Mgc-Map provided by the underlying joint distribution of the data. In other words, each relationship has a specific map that characterizes it based on its joint distribution, and Mgc is able to accurately estimate it via sample observations. The existence of a population level characterization of the joint distribution strongly differentiates Mgc from previously proposed multi-scale geometric or topological characterizations of data, such as persistence diagrams [23].

Mgc is Computationally Efficient Mgc does not incur large computational costs and has a similar complexity as existing methods. Though a naïve implementation of Mgc requires $\mathcal{O}(n^4)$ operations, we devised a nested implementation that requires only $\mathcal{O}(n^2 \log n)$ operations. Moreover, obtaining the Mgc-Map costs no additional computation, whereas other methods would require running a secondary computational step to decipher geometric properties of the relationship. Mgc can also trivially be parallelized, reducing computation to $\mathcal{O}(n^2 \log n/T)$, where T is the number of cores (see Algorithm C1 for details). Since T is often larger than $\log n$, in practice, Mgc can be $\mathcal{O}(n^2)$, meaning only a constant factor slower than Dcorr and Hsic, which is illustrated in Figure 6 of Shen et al. [30]. For example, at sample size $n = 5000$ and dimension $p = 1$, on a typical laptop computer, Dcorr requires around 0.5 seconds to compute the test statistic, whereas Mgc requires no more than 5 seconds. But the cost and time to obtain $2.5\times$ more data (so Dcorr has same average power as Mgc) typically far exceeds a few seconds. In comparison, the cost to compute a persistence diagram is typically $\mathcal{O}(n^3)$, which is orders of magnitude slower when $n > 10$. The running time of each method on the real data experiments are reported in Materials and Methods.

Mgc Uniquely Reveals Relationships in Real Data

Geometric intuition, numerical simulations, and theory all provide evidence that Mgc will be useful for real data discoveries. Nonetheless, real data applications provide another necessary ingredient to justify its use in practice. Below, we describe several real data applications where we have used Mgc to understand relationships in data that other methods were unable to provide.

Mgc Discovers the Relationships between Brain and Mental Properties

The human psyche is of course dependent on brain activity and structure. Previous work has studied two particular aspects of our psyche: personality and creativity, developing quantitative metrics for evaluating them using structured interviews [41, 42]. However, the relationship between brain activity and structure, and these aspects of our psyche, remains unclear [43–46]. For example, prior work did not evaluate the relationship between entire brain connectivity and all five factors of the standard personality model [41]. We therefore utilized Mgc to investigate published open access data (see Materials and Methods for details).

First, we analyzed the relationship between resting-state functional magnetic resonance (rs-fMRI) ac-

Table 2: The p-values for brain imaging vs mental properties. Mgc *always* uncovers the existence of significant relationships and discovers the underlying optimal scales. Bold indicates significant p-value per dataset.

Testing Pairs / Methods	Mgc	Dcorr	Mcorr	Hhg	Hsic
Activity vs Personality	0.043	0.667	0.441	0.059	0.124
Connectivity vs Creativity	0.011	0.010	0.011	0.031	0.092

tivity and personality [47]. The first row of Table 2 compares the p-value of different methods, and Figure 4A shows the Mgc-Map for the sample data. Mgc is able to yield a significant p-value (< 0.05), whereas all previously proposed global dependence tests under consideration (MANTEL, Dcorr, Mcorr, or HHG) fail to detect dependence at a significance level of 0.05. Moreover, the Mgc-Map provides a characterization of the dependence, for which the optimal scale indicates that the dependency is strongly nonlinear. Interestingly, the Mgc-Map does not look like any of the 20 images from the simulated data, suggesting that the nonlinearity characterizing this dependency is more complex or otherwise different from those we have considered so far.

Second, we investigated the relationship between diffusion MRI derived connectivity and creativity [42]. The second row of Table 2 shows that Mgc is able to ascertain a dependency between the whole brain network and the subject’s creativity. The Mgc-Map in Figure 4B closely resembles a linear relationship where the optimal scale is global. The close-to-linear relationship is also supported from the p-value table as all methods except Hsic are able to detect significant dependency, which suggests that there is relatively little to gain by pursuing nonlinear regression techniques, potentially saving valuable research time by avoiding tackling an unnecessary problem. The test statistic for both Mgc and Mcorr equal 0.04, which is quite close to zero despite a significant p-value, implying a relatively weak and noisy relationship. A prediction of creativity via linear regression turns out to be non-significant, which implies that the sample size is too low to obtain useful predictive accuracy (not shown), indicating that more data are required for single subject predictions. If one had first directly estimated the regression function, obtaining a null result, it would remain unclear whether a relationship existed. This experiment demonstrates that for high-dimensional and potentially structured data, Mgc is able to reveal dependency with relatively small sample size while parametric techniques and directly estimating regression functions can often be ineffective.

The performance in the real data closely matches the simulations results: the first dataset exhibits a strongly nonlinear relationship, for which Mgc has the lowest p-value, followed by HHG and Hsic and then all other methods; the second dataset exhibits a close-to-linear relationship, for which global methods perform the best while HHG and Hsic are trailing. Moreover, Mgc detected a complex nonlinear relationship for brain activity versus personality, and a nearly linear but noisy relationship for brain network versus creativity, the only method able to make either of those claims. In a separate experiment, we assessed the frequency with which Mgc obtained false positive results using brain activity data, based on experiments from [48, 49]. Supplementary Figure 5 shows that Mgc achieves a false positive rate of 5% when using a significance level of 0.05, implying that it correctly controls for false positives, unlike typical parametric methods on these data.

Mgc Identifies Potential Cancer Proteomics Biomarkers

Mgc can also be useful for a completely complementary set of scientific questions: screening proteomics data for biomarkers, often involving the analysis of tens of thousands of proteins, peptides, or transcripts in multiple samples representing a variety of disease types. Determining whether there is

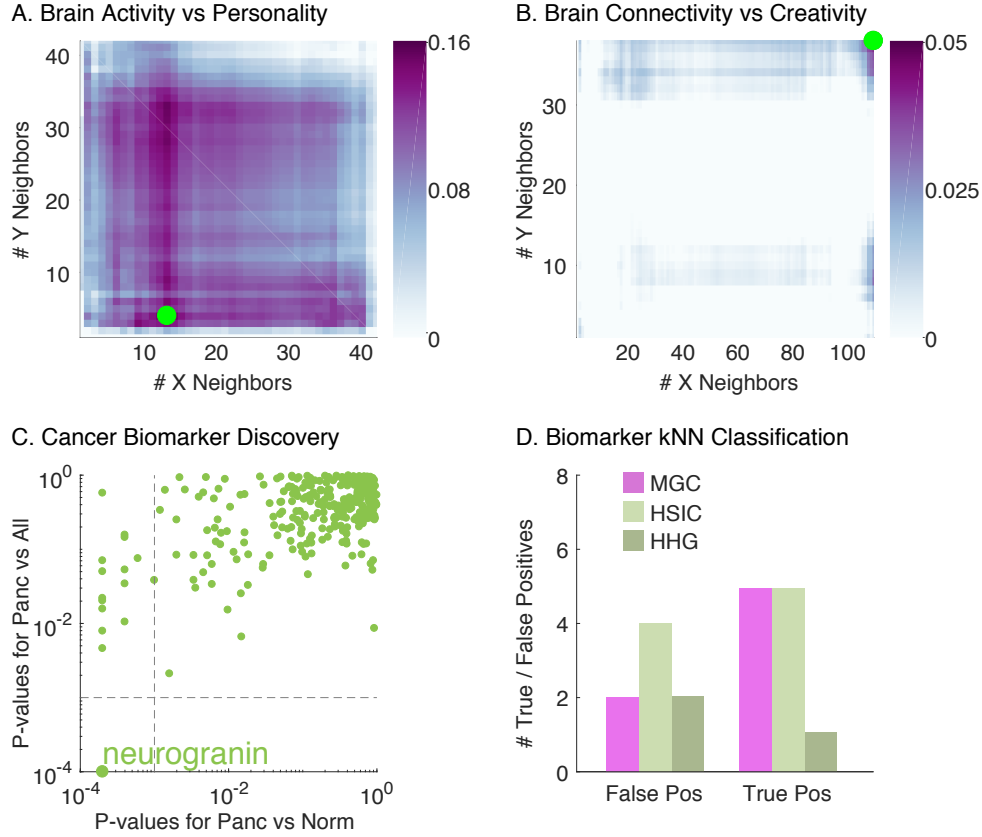


Figure 4: Demonstration that MGC successfully detects dependency, distinguishes linearity from nonlinearity, and identifies the most informative feature in a variety of real data experiments. **(A)** The MGC-Map for brain activity versus personality. MGC has a large test statistic and a significant p-value at the optimal scale (13, 4), while the global counterpart is non-significant. That the optimal scale is non-global implies a strongly nonlinear relationship. **(B)** The MGC-Map for brain connectivity versus creativity. The image is similar to that of a linear relationship, and the optimal scale equals the global scale, thus both MGC and MCOFF are significant in this case. **(C)** For each peptide, the x-axis shows the p-value for testing dependence between pancreatic and healthy subjects by MGC, and the y-axis shows the p-value for testing dependence between pancreatic and all other subjects by MGC. At critical level 0.05, MGC identifies a unique protein after multiple testing adjustment. **(D)** The true and false positive counts using a k-nearest neighbor (choosing the best $k \in [1, 10]$) leave-one-out classification using only the significant peptides identified by each testing method. The peptide identified by MGC achieves the best true and false positive rates, as compared to the peptides identified by HSIC or HHG.

a relationship between one or more of these markers and a particular disease state can be challenging, but is a necessary first step [50]. We sought to discover new useful protein biomarkers from a quantitative proteomics technique that measures protein and peptide abundance called Selected Reaction Monitoring (SRM) [51]. Specifically, we were interested in finding biomarkers that were unique to pancreatic cancer, because it is lethal and no clinically useful biomarkers are currently available [52].

The data consist of proteolytic peptides derived from the blood samples of 95 individuals harboring pancreatic ($n = 10$), ovarian ($n = 24$), colorectal cancer ($n = 28$), and healthy controls ($n = 33$). The processed data included 318 peptides derived from 121 proteins. Previously, we used these data and other techniques to find ovarian cancer biomarkers (a much easier task because the dataset has twice as many ovarian patients) and validated them with subsequent experiments [53]. Therefore, our first step was to check whether Mgc could correctly identify ovarian biomarkers. Indeed, the peptides that have been validated previously are also identified by Mgc. Emboldened, using the same dataset, we applied Mgc to screen for biomarkers unique to pancreatic cancer. To do so, we first screened for a difference between pancreatic cancer and healthy controls, identifying several potential biomarkers. Then, we screened for a difference between pancreatic cancer and all other conditions, to find peptides that differentiate pancreatic cancer from other cancers. Figure 4C shows the p-value of each peptide assigned by Mgc, which reveals one particular protein, neurogranin, that exhibits a strong dependency specifically with pancreatic cancer. Subsequent literature searches reveal that neurogranin is a potentially valuable biomarker for pancreatic cancer because it is exclusively expressed in brain tissue among normal tissues and has not been linked with any other cancer type [54, 55]. In comparison, Hsic identified neurogranin as well, but it also identified another peptide; Hhg identified the same two by Hsic, and a third peptide. A literature evaluation of these additional peptides shows that they are upregulated in other cancers as well and are unlikely to be useful as a pancreatic biomarker [56, 57]. The rest of the global methods did not identify any markers at significance level 0.05, see Materials and Methods for more details and Table 4 for identified peptide information using each method.

Since there is no ground truth yet in this experiment, we further carried out a classification task using the biomarkers identified by the various algorithms, using a k-nearest-neighbor classifier to predict pancreatic cancer, and a leave-one-subject-out validation. Figure 4D shows that the peptide selected by Mgc (neurogranin) works better than any other subset of the peptides selected by Hsic or Hhg, in terms of both fewer false positives and negatives. This analysis suggests Mgc can effectively be used for screening and subsequent classification.

Discussion

There are a number of connections between Mgc and other prominent statistical procedures that may be worth further exploration. First, Mgc can be thought of as a regularized or sparsified variant of distance or kernel methods. Regularization is central to high-dimensional and ill-posed problems, where dimensionality is larger than sample size. Second, Mgc can also be thought of as learning a metric because it chooses the optimal scale amongst a set of n^2 truncated distances, motivating studying the relationship between Mgc and recent advances in metric learning [58]. In particular, deep learning can be thought of as metric learning [59], and generative adversarial networks [60] are implicitly testing for equality, which is closely related to dependence [61]. While Mgc searches over a two-dimensional parameter space to optimize the metric, deep learning searches over a much larger parameter space, sometimes including millions of dimensions. Probably neither is optimal, and somewhere between the two would be useful in many tasks. Third, energy statistics provide state of the art approaches to other problems, including goodness-of-fit [62], analysis of variance [63], conditional dependence [64, 65], and feature selection [66, 67], so Mgc can be adapted for them as well. Indeed, Mgc can also implement

a two-sample (or generally the K-sample) test [19, 68, 69]. Specifically, for more than two modalities, one may use summation of pairwise Mgc test statistics, similar to how energy statistic is generalized to K-sample testing from two-sample testing [19, 63, 70], or how canonical correlation analysis is generalized into more than two modalities [71–73]. Finally, although energy statistics have not yet been explicitly used for classification, regression, or dimensionality reduction, Mgc opens the door to these applications by providing guidance as to how to proceed. Specifically, it is well documented in machine learning literature that the choice of kernel, metric, or scale often has a strong effect on the performance of different machine learning algorithms [28]. Mgc provides a mechanism to estimate scale that is both theoretically justified and computationally efficient, by optimizing a metric for a task wherein the previous methods lacked a notion of optimization. Nonlinear dimensionality reduction procedures, such as Isomap [74] and local linear embedding [75] for example, must also choose a scale, but have no principled criteria for doing so. Mgc could be used to provide insight into multimodal dimensionality reduction as well.

The default metric choice of Mgc in this paper is always the Euclidean distance, but other metric choices may be more appropriate in different fields, and using the strong negative type metric as specified in [15] guarantees consistency. However, if multiple metric choices are experimented to yield multiple Mgc p-values, then the optimal p-value should be properly corrected for multiple testing. Alternatively, one may use the maximum Mgc statistic among multiple metric choices, apply the same procedure in each permutation (i.e., in each permutation, use the same number of metric choices and take the maximum Mgc as the permuted statistic), then derive a single p-value. Such a testing procedure properly controls the type 1 error level without the need for additional correction.

Mgc also addresses a particularly vexing statistical problem that arises from the fact that methods for discovering dependencies are typically dissociated from methods for deciphering them. This dissociation creates a problem because the statistical assumptions underlying the “deciphering” methods become compromised in the process of “discovering”; this is called the “post-selection inference” problem [40]. The most straightforward way to address this issue is to collect new data, which is costly and time-consuming. Therefore, researchers often ignore this fact and make statistically invalid claims. Mgc circumvents this dilemma by carefully constructing its permutation test to estimate the scale in the process of estimating a p-value, rather than after. To our knowledge, Mgc is the first dependence test to take a step towards valid post-selection inference.

As a separate next theoretical extension, we could reduce the computational space and time required by Mgc. Mgc currently requires space and time that are quadratic with respect to the number of samples, which can be costly for very large data. Recent advances in related work demonstrated that one could reduce computational time of distance-based tests to close to linear via faster implementation, subsampling, random projection, and null distribution approximation [76–79], making it feasible for large amount of data. Alternately, semi-external memory implementations would allow running Mgc even as the interpoint comparison matrix exceeds the size of main memory [80–83].

Finally, Mgc is easy to use. Source code is available in MATLAB, R, and Python from <https://neurodata.io/tools/> and <https://github.com/neurodata/mgc>. Code for reproducing all the figures in this manuscript is also available from the above websites. We showed Mgc’s value in diverse applications spanning neuroscience (which motivated this work) and an ‘omics example. Applications in other domains facing similar questions of dependence, such as finance, pharmaceuticals, commerce, and security, could likewise benefit from Mgc.

Materials and Methods

Mathematical Details

This section contains essential mathematical details on independence testing, the notion of the generalized correlation coefficient and the distance-based correlation measure, how to compute the local correlations, and the smoothing technique. A statistical treatment on MGC is in [30], which introduces the population version of MGC and various theoretical properties.

Testing Independence

Given pairs of observations $(\mathbf{x}_i, \mathbf{y}_i) \in \mathbb{R}^p \times \mathbb{R}^q$ for $i = 1, \dots, n$, assume they are independently identically distributed as $(X, Y) \stackrel{iid}{\sim} F_{XY}$. If the two random variables X and Y are independent, the joint distribution equals the product of the marginals, i.e., $F_{XY} = F_X F_Y$. The statistical hypotheses for testing independence is as follows:

$$\begin{aligned} H_0 : F_{XY} &= F_X F_Y, \\ H_A : F_{XY} &\neq F_X F_Y. \end{aligned}$$

Given a test statistic, the testing power equals the probability of rejecting the independence hypothesis (i.e. the null hypothesis) when it is false. A test statistic is consistent if and only if the testing power increases to 1 as sample size increases to infinity. We would like a test to be universally consistent, i.e., consistent against all joint distributions. Dcorr, Mcorr, Hsic, and Hhg are all consistent against any joint distribution of finite second moments and finite dimension.

Note that p is the dimension for \mathbf{x} 's, q is the dimension for \mathbf{y} 's. For MGC and all benchmark methods, there is no restriction on the dimensions, i.e., the dimensions can be arbitrarily large, and p is not required to equal q . The ability to handle data of arbitrary dimension is crucial for modern big data. There also exist some special methods that only operate on one-dimensional data, such as [37, 69, 76], which are not directly applicable to multidimensional data.

Correlation Measures

To achieve consistent testing, most state-of-the-art dependence measures operate on pairwise comparisons, either similarities (such as kernels) or dissimilarities (such as distances).

Let $\mathcal{X}_n = \{\mathbf{x}_1, \dots, \mathbf{x}_n\} \in \mathbb{R}^{p \times n}$ and $\mathcal{Y}_n = \{\mathbf{y}_1, \dots, \mathbf{y}_n\} \in \mathbb{R}^{q \times n}$ denote the matrices of sample observations, and δ_x be the distance function for \mathbf{x} 's and δ_y for \mathbf{y} 's. One can then compute two $n \times n$ distance matrices $\tilde{A} = \{\tilde{a}_{ij}\}$ and $\tilde{B} = \{\tilde{b}_{ij}\}$, where $\tilde{a}_{ij} = \delta_x(\mathbf{x}_i, \mathbf{x}_j)$ and $\tilde{b}_{ij} = \delta_y(\mathbf{y}_i, \mathbf{y}_j)$. A common example of the distance function is the Euclidean metric (L^2 norm), which serves as the starting point for all methods in this manuscript.

Let A and B be the transformed (e.g., centered) versions of the distance matrices \tilde{A} and \tilde{B} , respectively. Any “generalized correlation coefficient” [35, 36] can be written as:

$$c(\mathcal{X}_n, \mathcal{Y}_n) = \frac{1}{z} \sum_{i=1}^n \sum_{j=1}^n a_{ij} b_{ij}, \quad (1)$$

where z is proportional to the standard deviations of A and B , that is $z = n^2 \sigma_a \sigma_b$. In words, c is the global sample correlation across *pairwise comparison matrices* A and B , and is normalized into the range $[-1, 1]$, which usually has expectation 0 under independence and implies a stronger dependency when the correlation is further away from 0.

Traditional correlations such as the Pearson's correlation and the rank correlation can be written via the above correlation formulation, by using A and B directly from sample observations rather than distances. Distance-based methods like **DCORR** and **MANTEL** operate on the Euclidean distance by default, or other metric choices on the basis of domain knowledge; then transform the resulting distance matrices \tilde{A} and \tilde{B} by certain centering schemes into A and B . **HSIC** chooses the Gaussian kernel and computes two kernel matrices, then transform the kernel matrices \tilde{A} and \tilde{B} by the same centering scheme as **DCORR**. For **MGC**, A and B are always distance matrices (or can be transformed to distances from kernels by [18]), and we shall apply a slightly different centering scheme that turns out to equal **DCORR**.

To carry out the hypothesis testing on sample data via a nonparametric test statistic, e.g., a generalized correlation, the permutation test is often an effective choice [84], because a p-value can be computed by comparing the correlation of the sample data to the correlation of the permuted sample data. The independence hypothesis is rejected if the p-value is lower than a pre-determined type 1 error level, say 0.05. Then the power of the test statistic equals the probability of a correct rejection at a specific type 1 error level. Note that **HNG** is the only exception that cannot be cast as a generalized correlation coefficient, but the permutation testing is similarly effective for the **HNG** test statistic; also note that the *iid* assumption is critical for permutation test to be valid, which may not be applicable in special cases like auto-correlated time series [85].

Distance Correlation (**DCORR**) and the Unbiased Version (**MCORR**)

Define the row and column means of \tilde{A} by $\bar{a}_{.j} = \frac{1}{n} \sum_{i=1}^n \tilde{a}_{ij}$ and $\bar{a}_{i.} = \frac{1}{n} \sum_{j=1}^n \tilde{a}_{ij}$. **DCORR** defines

$$a_{ij} = \begin{cases} \tilde{a}_{ij} - \bar{a}_{i.} - \bar{a}_{.j} + \bar{a}, & \text{if } i \neq j, \\ 0, & \text{if } i = j, \end{cases}$$

and similarly for b_{ij} . For distance correlation, the numerator of Equation 1 is named the distance covariance (**DCOV**), while σ_a and σ_b in the denominator are the square root of each distance variance. The centering scheme is important to guarantee the universal consistency of **DCORR**, whereas **MANTEL** uses a simple centering scheme and thus not universally consistent.

Let $c(X, Y)$ be the population distance correlation, that is, the distance correlation between the underlying random variables X and Y . Székely et al. (2007) [12] define the population distance correlation via the characteristic functions of F_X and F_Y , and show that the population distance correlation equals zero if and only if X and Y are independent, for any joint distribution F_{XY} of finite second moments and finite dimensionality. They also show that as $n \rightarrow \infty$, the sample distance correlation converges to the population distance correlation, that is, $c(\mathcal{X}_n, \mathcal{Y}_n) \rightarrow c(X, Y)$. Thus the sample distance correlation is consistent against any dependency of finite second moments and dimensionality. Of note, the distance covariance, distance variance, and distance correlation are always non-negative. Moreover, the consistency result holds for a much larger family of metrics, those of strong negative type [15].

It turns out that the sample distance correlation has a finite-sample bias, especially as the dimension p or q increases [14]. For example, for independent Gaussian distributions, the sample distance correlation converges to 1 as $p, q \rightarrow \infty$. By excluding the diagonal entries and slightly modifies the off-diagonal entries of \mathcal{A} and \mathcal{B} , Székely and Rizzo [14, 86] show that **MCORR** is an unbiased estimator of the population distance correlation $c(\mathbf{x}, \mathbf{y})$ for all p, q, n , which is approximately normal even if $p, q \rightarrow \infty$. Thus it enjoys the same theoretical consistency as **DCORR** and always has zero mean under independence.

Local Correlations

Given any matrices A and B , we can define a set of local correlations as follows. Let $R(A_{\cdot j}, i)$ be the “rank” of x_i relative to x_j , that is, $R(A_{\cdot j}, i) = k$ if x_i is the k^{th} closest point (or “neighbor”) to x_j , as determined by ranking the $n - 1$ distances to x_j . Define $R(B_{i \cdot}, j)$ equivalently for the Y ’s, but ranking relative to the rows rather than the columns (see below for explanation). For any neighborhood size k around each x_i and any neighborhood size l around each y_j , we define the local pairwise comparisons:

$$\tilde{a}_{ij}^k = \begin{cases} a_{ij}, & \text{if } R(A_{\cdot j}, i) \leq k, \\ 0, & \text{otherwise;} \end{cases} \quad \tilde{b}_{ij}^l = \begin{cases} b_{ij}, & \text{if } R(B_{i \cdot}, j) \leq l, \\ 0, & \text{otherwise;} \end{cases} \quad (2)$$

and then let $a_{ij}^k = \tilde{a}_{ij}^k - \bar{a}^k$, where \bar{a}^k is the mean of $\{\tilde{a}_{ij}^k\}$, and similarly for b_{ij}^l .

The *local* correlation coefficient at a given scale is defined to effectively exclude large distances:

$$c^{kl}(\mathcal{X}_n, \mathcal{Y}_n) = \frac{1}{z_{kl}} \sum_{i,j=1}^n a_{ij}^k b_{ij}^l, \quad (3)$$

where $z_{kl} = n^2 \sigma_a^k \sigma_b^l$, with σ_a^k and σ_b^l is the standard deviations for the truncated pairwise comparisons. The MGC-Map can be constructed by computing all local correlations, which allows the discovery of the optimal correlation. For any aforementioned correlation (DCORR, MCORR, HSIC, MANTEL, PEARSON), one can define its local correlations by using Equation 3 and plugging in the respective a_{ij} and b_{ij} from Equation 1.

As most nonlinear relationships intrinsically exhibit a local linear structure, considering the nearest-neighbors is able to amplify the dependency signal over the global correlation. There could be two other scenarios: when the small distances in one modality mostly correspond to large distances in another modality, or when the large distances in one modality correspond to large distance in another modality. For the first scenario, the small distances become negative terms after centering while the large distances become positive terms after centering, so adding their product to c^{kl} will cause the test statistic to be smaller — in fact, as distance correlation is shown to be > 0 under dependence [12], the first scenario cannot happen for all distances pairs. For the second scenario, one can experiment using the large distances (or the furthest neighbors) only by reversing the ranking scheme in local correlation to descending order. However, whenever the large distances are highly correlated, the small distances must also be highly correlated after centering by the mean distances, so global correlation coefficient like DCORR already handles this scenario. Therefore considering the nearest-neighbor may significantly improve the performance over global correlation, while considering the other scenarios does not.

MGC as the Optimal Local Correlation

We define the multiscale graph correlation statistic as the optimal local correlation, for which the family of local correlation is computed based on Euclidean distance and MCORR transformation.

Instead of taking a direct maximum, MGC takes a smoothed maximum, i.e., the maximum local correlation of the largest connected component R such that all local correlations within R are significant. If no such region exists, MGC defaults the test statistic to the global correlation (details in Algorithm C2). Thus, we can write:

$$c^*(\mathcal{X}_n, \mathcal{Y}_n) = \max_{(k,l) \in R} c^{kl}(\mathcal{X}_n, \mathcal{Y}_n) \quad (4)$$

$$R = \text{Largest Connected Component of } \{(k, l) \text{ such that } c^{kl} > \max(\tau, c^{nn})\}.$$

Then the optimal scale equals all scales within R whose local correlations are as large as c^* . The choice of τ is made explicit in the pseudo-code, with further discussion and justification offered in [30].

Proof for Theorem 1

Theorem 1. *When (X, Y) are linearly related (rotation, scaling, translation, reflection), the optimal scale of Mgc equals the global scale. Conversely, that the optimal scale is local implies a nonlinear relationship.*

Proof. It suffices to prove the first statement, then the second statement follows by contrapositive. When (X, Y) are linearly related, $Y = WX + b$ for a unitary matrix W and a constant b up-to possible scaling, in which case the distances are preserved, i.e., $\|y_i - y_j\| = \|Wx_i - Wx_j\| = \|x_i - x_j\|$. It follows that $\text{Mcorr}(\mathcal{X}_n, \mathcal{Y}_n) = 1$, so the global scale achieves the maximum possible correlation, and the largest connected region R is empty. Thus the optimal scale is global and $\text{Mgc}(\mathcal{X}_n, \mathcal{Y}_n) = \text{Mcorr}(\mathcal{X}_n, \mathcal{Y}_n) = 1$. \square

Computational Complexity of Each Step

The distance computation takes $\mathcal{O}(n^2 \max\{p, q\})$, and the ranking process takes $\mathcal{O}(n^2 \log n)$. Once the distance and ranking are completed, computing one local generalized correlation requires $\mathcal{O}(n^2)$ (see Algorithm C4). Thus a naive approach to compute all local generalized correlations requires at least $\mathcal{O}(n^2 \max\{n^2, p, q\})$ by going through all possible scales, meaning possibly $\mathcal{O}(n^4)$ which would be computationally prohibitive. However, given the distance and ranking information, we devised an algorithm that iteratively computes all local correlations in $\mathcal{O}(n^2)$ by re-using adjacent smaller local generalized correlations (see Algorithm C5). Therefore, when including the distance computation and ranking overheads, the MGC statistic is computed in $\mathcal{O}(n^2 \max\{\log n, p, q\})$, which has the same running time as the HHG statistic, and the same running time up to a factor of $\log n$ as global correlations like Dcorr and Mcorr, which require $\mathcal{O}(n^2 \max\{p, q\})$ time. By utilizing a multi-core architecture, MGC can be computed in $\mathcal{O}(n^2 \max\{\log n, p, q\}/T)$ instead. As $T = \log(n)$ is often a small number, e.g., T is no more than 30 at 1 billion samples, thus MGC can be effectively computed in the same complexity as Dcorr. Note that the permutation test adds another r random permutations to the n^2 term, so computing the p-value requires $\mathcal{O}(n^2 \max\{\log n, p, q, r\}/T)$.

MGC Algorithms and Testing Procedures

Six algorithms are presented in order:

- Algorithm C1 describes MGC in its entirety (which calls most of the other algorithms as functions).
- Algorithm C2 computes the MGC test statistic.
- Algorithm C3 computes the p-value of MGC by the permutation test.
- Algorithm C4 computes the local generalized correlation coefficient at a given scale (k, l) , for a given choice of the global correlation coefficient.
- Algorithm C5 efficiently computes all local generalized correlations, in nearly the same running time complexity as computing one local generalized correlation.
- Algorithm C6 evaluates the testing power of MGC by a given distribution.

For ease of presentation, we assume there are no repeating observations of X or Y , and note that Mcorr is the global correlation choice that MGC builds on.

Pseudocode C1 Multiscale Graph Correlation (MGC); requires $\mathcal{O}(n^2 \max(\log n, p, q, r)/T)$ time, where r is the number of permutations and T is the number of cores available for parallelization.

Input: n samples of (x_i, y_i) pairs, an integer r for the number of random permutations.

Output: (i) MGC statistic c^* , (ii) the optimal scale (k, l) , (iii) the p-value $p(c^*)$,

function MGC((x_i, y_i) , for $i \in [n]$)

(1) Calculate all pairwise distances:

for $i, j := 1, \dots, n$ **do**

$a_{ij} = \delta_x(x_i, x_j)$

$\triangleright \delta_x$ is the distance between pairs of x samples

$b_{ij} = \delta_y(y_i, y_j)$

$\triangleright \delta_y$ is the distance between pairs of y samples

end for

 Let $A = \{a_{ij}\}$ and $B = \{b_{ij}\}$.

(2) Calculate Multiscale Correlation Map \mathcal{C} & MGC Test Statistic:

$[c^*, \mathcal{C}, k, l] = \text{MGCSAMPLESTAT}(A, B)$

\triangleright Algorithm C2

(3) Calculate the p-value

$pval(c^*) = \text{PERMUTATIONTEST}(A, B, r, c^*)$

\triangleright Algorithm C3

end function

Pseudocode C2 Mgc test statistic. This algorithm computes all local correlations, take the smoothed maximum, and reports the (k, l) pair that achieves it. For the smoothing step, it: (i) finds the largest connected region in the correlation map, such that each correlation is significant, i.e., larger than a certain threshold to avoid correlation inflation by sample noise, (ii) take the largest correlation in the region, (iii) if the region area is too small, or the smoothed maximum is no larger than the global correlation, the global correlation is used instead. The running time is $\mathcal{O}(n^2)$.

Input: A pair of distance matrices $(A, B) \in \mathbb{R}^{n \times n} \times \mathbb{R}^{n \times n}$.

Output: The Mgc statistic $c^* \in \mathbb{R}$, all local statistics $\mathcal{C} \in \mathbb{R}^{n \times n}$, and the corresponding local scale $(k, l) \in \mathbb{N} \times \mathbb{N}$.

```

1: function MGCSAMPLESTAT( $A, B$ )
2:    $\mathcal{C} = \text{MGCALLLOCAL}(A, B)$  ▷ All local correlations
3:    $\tau = \text{THRESHOLDING}(\mathcal{C})$  ▷ find a threshold to determine large local correlations
4:   for  $i, j := 1, \dots, n$  do  $r_{ij} \leftarrow \mathbb{I}(c^{ij} > \tau)$  end for ▷ identify all scales with large correlation
5:    $\mathcal{R} \leftarrow \{r_{ij} : i, j = 1, \dots, n\}$  ▷ binary map encoding scales with large correlation
6:    $\mathcal{R} = \text{CONNECTED}(\mathcal{R})$  ▷ largest connected component of the binary matrix
7:    $c^* \leftarrow \mathcal{C}(n, n)$  ▷ use the global correlation by default
8:    $k \leftarrow n, l \leftarrow n$ 
9:   if  $(\sum_{i,j} r_{ij}) \geq 2n$  then ▷ proceed when the significant region is sufficiently large
10:     $[c^*, k, l] \leftarrow \max(\mathcal{C} \circ \mathcal{R})$  ▷ find the smoothed maximum and the respective scale
11:   end if
12: end function

```

Input: $\mathcal{C} \in \mathbb{R}^{n \times n}$.

Output: A threshold τ to identify large correlations.

```

13: function THRESHOLDING( $\mathcal{C}$ )
14:    $\tau \leftarrow \sum_{c^{ij} < 0} (c^{ij})^2 / \sum_{c^{ij} < 0} 1$  ▷ variance of all negative local generalized correlations
15:    $\tau \leftarrow \max\{0.01, \sqrt{\tau}\} \times 3.5$  ▷ threshold based on negative correlations
16:    $\tau \leftarrow \max\{\tau, 2/n, c^{nn}\}$ 
17: end function

```

Pseudocode C3 Permutation Test. This algorithm uses the random permutation test with r random permutations for the p-value, requiring $\mathcal{O}(rn^2 \log n)$ for MGC. In the real data experiment we always set $r = 10,000$. Note that the p-value computation for any other global generalized correlation coefficient follows from the same algorithm by replacing MGC with the respective test statistic.

Input: A pair of distance matrices $(A, B) \in \mathbb{R}^{n \times n} \times \mathbb{R}^{n \times n}$, the number of permutations r , and MGC statistic c^* for the observed data.

Output: The p-value $pval \in [0, 1]$.

```

1: function PERMUTATIONTEST( $A, B, r, c^*$ )
2:   for  $t := 1, \dots, r$  do
3:      $\pi = \text{RANDPERM}(n)$                                 ▷ generate a random permutation of size  $n$ 
4:      $c_0^*[t] = \text{MGCSAMPLESTAT}(A, B(\pi, \pi))$              ▷ calculate the permuted MGC statistic
5:   end for
6:    $pval(c^*) \leftarrow \frac{1}{r} \sum_{t=1}^r \mathbf{I}(c^* \leq c_0^*[t])$        ▷ compute p-value of MGC
7: end function

```

Pseudocode C4 Compute local test statistic at a given scale. This algorithm runs in $\mathcal{O}(n^2)$ once the rank information is provided, which is suitable for MGC computation if an optimal scale is already estimated. But it would take $\mathcal{O}(n^4)$ if used to compute all local generalized correlations. Note that for the default MGC implementation uses single centering, the centering function centers A by column and B by row, and the sorting function sorts A within column and B within row. By utilizing $T = \log(n)$ cores, the sorting function can be easily parallelized to take $\mathcal{O}(n^2 \log(n)/T) = \mathcal{O}(n^2)$.

Input: A pair of distance matrices $(A, B) \in \mathbb{R}^{n \times n} \times \mathbb{R}^{n \times n}$, and a local scale $(k, l) \in \mathbb{N} \times \mathbb{N}$.

Output: The local generalized correlation coefficient $c^{kl} \in [-1, 1]$.

```

1: function LOCALGENCORR( $A, B, k, l$ )
2:   for  $Z := A, B$  do  $\mathcal{E}^Z = \text{SORT}(Z)$  end for                                ▷ parallelized sorting
3:   for  $Z := A, B$  do  $Z = \text{CENTER}(Z)$  end for                                ▷ center distance matrices
4:    $\tilde{c}^{kl} \leftarrow \text{tr}((A \circ \mathcal{E}^A)^\top \times (B \circ (\mathcal{E}^B)^\top))$                     ▷ un-normalized local distance covariance
5:    $v^A \leftarrow \text{tr}((A \circ \mathcal{E}^A)^\top \times (A \circ (\mathcal{E}^A)^\top))$                         ▷ local distance variances
6:    $v^B \leftarrow \text{tr}((B \circ \mathcal{E}^B)^\top \times (B \circ (\mathcal{E}^B)^\top))$ 
7:    $e^A \leftarrow \sum_{i,j=1}^n (A \circ \mathcal{E}^A)_{ij}$                                         ▷ sample means
8:    $e^B \leftarrow \sum_{i,j=1}^n (B \circ \mathcal{E}^B)_{ij}$ 
9:    $c^{kl} \leftarrow (\tilde{c}^{kl} - e^A e^B / n^2) / \sqrt{(v^A - (e^A/n)^2)(v^B - (e^B/n)^2)}$     ▷ center and normalize
10: end function

```

Pseudocode C5 Compute the multiscale correlation map (i.e., all local generalized correlations) in $\mathcal{O}(n^2 \log n/T)$. Once the distances are sorted, the remaining algorithm runs in $\mathcal{O}(n^2)$. An important observation is that each product $a_{ij}b_{ij}$ is included in c^{kl} if and only if (k, l) satisfies $k \leq R(A_{\cdot j}, i)$ and $l \leq R(B_{\cdot j}, i)$, so it suffices to iterate through $a_{ij}b_{ij}$ for $i, j := 1, \dots, n$, and add the product simultaneously to all c^{kl} whose scales are no more than $(R(A_{\cdot j}, i), R(B_{\cdot j}, i))$. To achieve the above, we iterate through each product, add it to c^{kl} at $(kl) = (R(A_{\cdot j}, i), R(B_{\cdot j}, i))$ only (so only one local scale is accessed for each operation); then add up adjacent c^{kl} for $k, l = 1, \dots, n$. The same applies to all local covariances, variances, and expectations.

Input: A pair of distance matrices $(A, B) \in \mathbb{R}^{n \times n} \times \mathbb{R}^{n \times n}$.

Output: The multiscale correlation map $\mathcal{C} \in [-1, 1]^{n \times n}$ for $k, l = 1, \dots, n$.

```

1: function MGCALLLOCAL( $A, B$ )
2:   for  $Z := A, B$  do  $\mathcal{E}^Z = \text{SORT}(Z)$  end for
3:   for  $Z := A, B$  do  $Z = \text{CENTER}(Z)$  end for
4:   for  $i, j := 1, \dots, n$  do                                ▷ iterate through all local scales to calculate each term
5:      $k \leftarrow \mathcal{E}_{ij}^Z$ 
6:      $l \leftarrow \mathcal{E}_{ij}^Z$ 
7:      $\tilde{c}^{kl} \leftarrow \tilde{c}^{kl} + a_{ij}b_{ij}$ 
8:      $v_k^A \leftarrow v_k^A + a_{ij}^2$ 
9:      $v_l^B \leftarrow v_l^B + b_{ij}^2$ 
10:     $e_k^A \leftarrow e_k^A + a_{ij}$ 
11:     $e_l^B \leftarrow e_l^B + b_{ij}$ 
12:   end for
13:   for  $k := 1, \dots, n-1$  do                                ▷ iterate through each scale again and add up adjacent terms
14:      $\tilde{c}^{1,k+1} \leftarrow \tilde{c}^{1,k} + \tilde{c}^{1,k+1}$ 
15:      $\tilde{c}^{k+1,1} \leftarrow \tilde{c}^{k+1,1} + \tilde{c}^{k+1,1}$ 
16:     for  $Z := A, B$  do  $v_{k+1}^Z \leftarrow v_k^Z + v_{k+1}^Z$  end for
17:     for  $Z := A, B$  do  $e_{k+1}^Z \leftarrow e_k^Z + e_{k+1}^Z$  end for
18:   end for
19:   for  $k, l := 1, \dots, n-1$  do
20:      $\tilde{c}^{k+1,l+1} \leftarrow \tilde{c}^{k+1,l} + \tilde{c}^{k,l+1} + \tilde{c}^{k+1,l+1} - \tilde{c}^{k,l}$ 
21:   end for
22:   for  $k, l := 1, \dots, n$  do
23:      $c^{kl} \leftarrow (\tilde{c}^{kl} - e_k^A e_l^B / n^2) / \sqrt{(v_k^A - e_k^{A^2} / n^2)(v_l^B - e_l^{B^2} / n^2)}$ 
24:   end for
25: end function

```

Pseudocode C6 Power computation of MGC against a given distribution. By repeatedly sampling from the joint distribution F_{XY} , sample data of size n under the null and the alternative are generated for r Monte-Carlo replicates. The power of MGC follows by computing the test statistic under the null and the alternative using Algorithm C2. In the simulations we use $r = 10,000$ MC replicates. Note that power computation for other benchmarks follows from the same algorithm by plugging in the respective test statistic.

Input: A joint distribution F_{XY} , the sample size n , the number of MC replicates r , and the type 1 error level α .

Output: The power β of MGC.

```

1: function MGCPOWER( $F_{XY}, n, r, \alpha$ )
2:   for  $t := 1, \dots, r$  do
3:     for  $i := [n]$  do
4:        $x_i^0 \stackrel{iid}{\sim} F_X, y_i^0 \stackrel{iid}{\sim} F_Y$  ▷ sample from null
5:        $(x_i^1, y_i^1) \stackrel{iid}{\sim} F_{XY}$ , ▷ sample from alternative
6:     end for
7:     for  $i, j := 1, \dots, n$  do
8:        $a_{ij}^0 = \delta_x(x_i^0, x_j^0), b_{ij}^0 = \delta_y(y_i^0, y_j^0)$  ▷ pairwise distances under the null
9:        $a_{ij}^1 = \delta_x(x_i^1, x_j^1), b_{ij}^1 = \delta_y(y_i^1, y_j^1)$  ▷ pairwise distances under the alternative
10:    end for
11:     $c_0^*[t] = \text{MGCSAMPLESTAT}(A^0, B^0)$  ▷ MGC statistic under the null
12:     $c_1^*[t] = \text{MGCSAMPLESTAT}(A^1, B^1)$  ▷ MGC statistic under the alternative
13:  end for
14:   $\omega_\alpha \leftarrow \text{CDF}_{1-\alpha}(c_0^*[t], t \in [r])$  ▷ the critical value of MGC under the null
15:   $\beta \leftarrow \sum_{t=1}^r (c_1^*[t] > \omega_\alpha) / r$  ▷ compute power by the alternative distribution
16: end function

```

Simulation Dependence Functions

This section provides the 20 different dependency functions used in the simulations. We used essentially the exact same relationships as previous publications to ensure a fair comparison [12, 38, 39]. We only made changes to add white noise and a weight vector for higher dimensions, thereby making them more difficult, to better compare all methods throughout different dimensions and sample sizes. A few additional relationships are also included.

For each sample $\mathbf{x} \in \mathbb{R}^p$, we denote $\mathbf{x}_{[d]}, d = 1, \dots, p$ as the d^{th} dimension of the vector \mathbf{x} . For the purpose of high-dimensional simulations, $\mathbf{w} \in \mathbb{R}^p$ is a decaying vector with $w_{[d]} = 1/d$ for each d , such that $\mathbf{w}^\top \mathbf{x}$ is a weighted summation of all dimensions of \mathbf{x} . Furthermore, $\mathcal{U}(a, b)$ denotes the uniform distribution on the interval (a, b) , $\mathcal{B}(p)$ denotes the Bernoulli distribution with probability p , $\mathcal{N}(\mu, \Sigma)$ denotes the normal distribution with mean μ and covariance Σ , U and V represent some auxiliary random variables, κ is a scalar constant to control the noise level (which equals 1 for one-dimensional simulations and 0 otherwise), and ϵ is a white noise from independent standard normal distribution unless mentioned otherwise.

For all of the below equations, $(X, Y) \stackrel{iid}{\sim} F_{XY} = F_{Y|X}F_X$. For each relationship, we provide the space of (X, Y) , and define $F_{Y|X}$ and F_X , as well as any additional auxiliary distributions.

1. Linear $(X, Y) \in \mathbb{R}^p \times \mathbb{R}$,

$$\begin{aligned} X &\sim \mathcal{U}(-1, 1)^p, \\ Y &= \mathbf{w}^\top X + \kappa\epsilon. \end{aligned}$$

2. Exponential $(X, Y) \in \mathbb{R}^p \times \mathbb{R}$:

$$\begin{aligned} X &\sim \mathcal{U}(0, 3)^p, \\ Y &= \exp(\mathbf{w}^\top X) + 10\kappa\epsilon. \end{aligned}$$

3. Cubic $(X, Y) \in \mathbb{R}^p \times \mathbb{R}$:

$$\begin{aligned} X &\sim \mathcal{U}(-1, 1)^p, \\ Y &= 128(\mathbf{w}^\top X - \frac{1}{3})^3 + 48(\mathbf{w}^\top X - \frac{1}{3})^2 - 12(\mathbf{w}^\top X - \frac{1}{3}) + 80\kappa\epsilon. \end{aligned}$$

4. Joint normal $(X, Y) \in \mathbb{R}^p \times \mathbb{R}^p$: Let $\rho = 1/2p$, I_p be the identity matrix of size $p \times p$, J_p be the matrix of ones of size $p \times p$, and $\Sigma = \begin{bmatrix} I_p & \rho J_p \\ \rho J_p & (1 + 0.5\kappa)I_p \end{bmatrix}$. Then

$$(X, Y) \sim \mathcal{N}(0, \Sigma).$$

5. Step Function $(X, Y) \in \mathbb{R}^p \times \mathbb{R}$:

$$\begin{aligned} X &\sim \mathcal{U}(-1, 1)^p, \\ Y &= \mathbf{I}(\mathbf{w}^\top X > 0) + \epsilon, \end{aligned}$$

where \mathbf{I} is the indicator function, that is $\mathbf{I}(z)$ is unity whenever z true, and zero otherwise.

6. Quadratic $(X, Y) \in \mathbb{R}^p \times \mathbb{R}$:

$$\begin{aligned} X &\sim \mathcal{U}(-1, 1)^p, \\ Y &= (\mathbf{w}^\top X)^2 + 0.5\kappa\epsilon. \end{aligned}$$

7. **W Shape** $(X, Y) \in \mathbb{R}^p \times \mathbb{R}$: $U \sim \mathcal{U}(-1, 1)^p$,

$$X \sim \mathcal{U}(-1, 1)^p,$$

$$Y = 4 \left[\left((w^\top X)^2 - \frac{1}{2} \right)^2 + w^\top U / 500 \right] + 0.5\kappa\epsilon.$$

8. **Spiral** $(X, Y) \in \mathbb{R}^p \times \mathbb{R}$: $U \sim \mathcal{U}(0, 5)$, $\epsilon \sim \mathcal{N}(0, 1)$,

$$X_{[d]} = U \sin(\pi U) \cos^d(\pi U) \text{ for } d = 1, \dots, p-1,$$

$$X_{[p]} = U \cos^p(\pi U),$$

$$Y = U \sin(\pi U) + 0.4p\epsilon.$$

9. **Uncorrelated Bernoulli** $(X, Y) \in \mathbb{R}^p \times \mathbb{R}$: $U \sim \mathcal{B}(0.5)$, $\epsilon_1 \sim \mathcal{N}(0, I_p)$, $\epsilon_2 \sim \mathcal{N}(0, 1)$,

$$X \sim \mathcal{B}(0.5)^p + 0.5\epsilon_1,$$

$$Y = (2U - 1)w^\top X + 0.5\epsilon_2.$$

10. **Logarithmic** $(X, Y) \in \mathbb{R}^p \times \mathbb{R}^p$: $\epsilon \sim \mathcal{N}(0, I_p)$

$$X \sim \mathcal{N}(0, I_p),$$

$$Y_{[d]} = 2 \log_2(|X_{[d]}|) + 3\kappa\epsilon_{[d]} \text{ for } d = 1, \dots, p.$$

11. **Fourth Root** $(X, Y) \in \mathbb{R}^p \times \mathbb{R}$:

$$X \sim \mathcal{U}(-1, 1)^p,$$

$$Y = |w^\top X|^{\frac{1}{4}} + \frac{\kappa}{4}\epsilon.$$

12. **Sine Period 4π** $(X, Y) \in \mathbb{R}^p \times \mathbb{R}^p$: $U \sim \mathcal{U}(-1, 1)$, $V \sim \mathcal{N}(0, 1)^p$, $\theta = 4\pi$,

$$X_{[d]} = U + 0.02pV_{[d]} \text{ for } d = 1, \dots, p,$$

$$Y = \sin(\theta X) + \kappa\epsilon.$$

13. **Sine Period 16π** $(X, Y) \in \mathbb{R}^p \times \mathbb{R}^p$: Same as above except $\theta = 16\pi$ and the noise on Y is changed to $0.5\kappa\epsilon$.

14. **Square** $(X, Y) \in \mathbb{R}^p \times \mathbb{R}^p$: Let $U \sim \mathcal{U}(-1, 1)$, $V \sim \mathcal{U}(-1, 1)$, $\epsilon \sim \mathcal{N}(0, 1)^p$, $\theta = -\frac{\pi}{8}$. Then

$$X_{[d]} = U \cos \theta + V \sin \theta + 0.05p\epsilon_{[d]},$$

$$Y_{[d]} = -U \sin \theta + V \cos \theta,$$

for $d = 1, \dots, p$.

15. **Two Parabolas** $(X, Y) \in \mathbb{R}^p \times \mathbb{R}$: $\epsilon \sim \mathcal{U}(0, 1)$, $U \sim \mathcal{B}(0.5)$,

$$X \sim \mathcal{U}(-1, 1)^p,$$

$$Y = \left((w^\top X)^2 + 2\kappa\epsilon \right) \cdot \left(U - \frac{1}{2} \right).$$

16. **Circle** $(X, Y) \in \mathbb{R}^p \times \mathbb{R}$: $U \sim \mathcal{U}(-1, 1)^p$, $\epsilon \sim \mathcal{N}(0, I_p)$, $r = 1$,

$$X_{[d]} = r \left(\sin(\pi U_{[d+1]}) \prod_{j=1}^d \cos(\pi U_{[j]}) + 0.4\epsilon_{[d]} \right) \text{ for } d = 1, \dots, p-1,$$

$$X_{[p]} = r \left(\prod_{j=1}^p \cos(\pi U_{[j]}) + 0.4\epsilon_{[p]} \right),$$

$$Y = \sin(\pi U_{[1]}).$$

17. Ellipse $(X, Y) \in \mathbb{R}^p \times \mathbb{R}$: Same as above except $r = 5$.
18. Diamond $(X, Y) \in \mathbb{R}^p \times \mathbb{R}^p$: Same as “Square” except $\theta = -\frac{\pi}{4}$.
19. Multiplicative Noise $(\mathbf{x}, \mathbf{y}) \in \mathbb{R}^p \times \mathbb{R}^p$: $u \sim \mathcal{N}(0, I_p)$,

$$\begin{aligned}\mathbf{x} &\sim \mathcal{N}(0, I_p), \\ \mathbf{y}_{[d]} &= u_{[d]} \mathbf{x}_{[d]} \text{ for } d = 1, \dots, p.\end{aligned}$$

20. Multimodal Independence $(X, Y) \in \mathbb{R}^p \times \mathbb{R}^p$: Let $U \sim \mathcal{N}(0, I_p)$, $V \sim \mathcal{N}(0, I_p)$, $U' \sim \mathcal{B}(0.5)^p$, $V' \sim \mathcal{B}(0.5)^p$. Then

$$\begin{aligned}X &= U/3 + 2U' - 1, \\ Y &= V/3 + 2V' - 1.\end{aligned}$$

For each distribution, X and Y are dependent except (20); for some relationships (8,14,16-18) they are independent upon conditioning on the respective auxiliary variables, while for others they are “directly” dependent. A visualization of each dependency with $D = D_y = 1$ is shown in Figure 2-Figure Supplement 1.

For the increasing dimension simulation in the main paper, we always set $\kappa = 0$ and $n = 100$, with p increasing. Note that $q = p$ for types 4, 10, 12, 13, 14, 18, 19, 20, otherwise $q = 1$. The decaying vector w is utilized for $p > 1$ to make the high-dimensional relationships more difficult (otherwise, additional dimensions only add more signal). For the one-dimensional simulations, we always set $p = q = 1$, $\kappa = 1$ and $n = 100$.

Real Data Processing

Brain Activity vs Personality

This experiment investigates whether there is any dependency between resting brain activity and personality. Human personality has been intensively studied for many decades; the most widely used and studied approach is the NEO Personality Inventory-Revised the characterized personality along five dimensions [41]. This dataset consists of 42 subjects, each with 197 time-steps of resting-state functional magnetic resonance activity (rs-fMRI) activity, as well as the subject’s five-dimensional “personality”. Adelstein et al. [47] were able to detect dependence between the activity of certain brain regions and dimensions of personality, but lacked the tools to test for dependence of whole brain activity against all five dimensions of personality. For the five-factor personality modality, we used the Euclidean distance. For the brain activity modality, we derived the following comparison function. For each scan, (i) run Configurable Pipeline for the Analysis of Connectomes pipeline [87] to process the raw brain images yielding a parcellation into 197 regions of interest, (ii) run a spectral analysis on each region and keep the power of band, (iii) bandpass and normalize it to sum to one, (iv) calculate the Kullback-Leibler divergence across regions to obtain a similarity matrix across comparing all regions. Then, use the normalized Hellinger distance to compute distances between each subject.

Brain Connectivity vs Creativity

This experiment investigates whether there is any dependency between brain structural networks and creativity. Creativity has been extensively studied in psychology; the “creativity composite index” (CCI) is an index similar to an “intelligence quotient” but for creativity rather than intelligence [42]. This dataset consists of 109 subjects, each with diffusion weighted MRI data as well as the subject’s CCI. Neural correlates of CCI have previously been investigated, though largely using structural MRI and cortical thickness [42]. Previously published results explored the relationship between graphs and CCI [88], but did not provide a valid test. We used Euclidean distance to compare CCI values. For the raw brain imaging data, we derived the following comparison function. For each scan we estimated brain networks from diffusion and structural MRI data via MIGRAINE, a pipeline for estimating brain networks from diffusion data [89]. We compute the distance between the graphs using the semi-parametric graph test statistic [90–92], embedding each graph into two dimensions and aligning the embeddings via a Procrustes analysis.

Proteins vs Cancer

This experiment investigated whether there is any dependency between abundance levels of peptides in human plasma and the presence of cancers. Selected Reaction Monitoring (SRM) is a targeted quantitative proteomics technique for measuring protein and peptide abundance in complicated biological samples [51]. In a previous study, we used SRM to identify 318 peptides from 33 normal, 10 pancreatic cancer, 28 colorectal cancer, and 24 ovarian cancer samples [53]. Then, using other methods, we identified three peptides that were implicated in ovarian cancer, and validated them as legitimate biomarkers with a follow-up experiment.

In this study, we performed the following five sets of tests on those data:

1. ovarian vs. normal for all proteins,
2. ovarian vs. normal for each individual protein,
3. pancreas vs. normal for all proteins,

4. pancreas vs. all others for each individual protein,
5. pancreas vs. normal for each individual protein.

These tests are designed to first validate the Mgc method from ovarian cancer, then identify biomarkers unique to pancreatic cancer, that is, find a protein that is able to tell the difference between pancreas and normals, as well as pancreas vs all other cancers. For each of the five tests, we create a binary label vector, with 1 indicating the cancer type of interest for the corresponding subject, and 0 otherwise. Then each algorithm is applied to each task. For all tests we used Euclidean distances and the type 1 error level is set to $\alpha = 0.05$. The three test sets assessing individual proteins provide 318 p-values; we used the Benjamini-Hochberg procedure [93] to control the false discovery rate. A summary of the results are reported in Table 3.

Table 3: Results for cancer peptide screening. The first two rows report the p-values for the tests of interest based on all peptides. The next four rows report the number of significant proteins from individual peptide tests; the Benjamini-Hochberg procedure is used to locate the significant peptides by controlling the false discovery rate at 0.05.

	Testing Pairs / Methods	Mgc	MANTEL	DCORR	MCORR	Hhg
1	Ovar vs. Norm: p-value	0.0001	0.0001	0.0001	0.0001	0.0001
2	Ovar vs. Norm: # peptides	218	190	186	178	225
3	Pancr vs. Norm: p-value	0.0082	0.0685	0.0669	0.0192	0.0328
4	Panc vs. Norm: # peptides	9	7	6	7	11
5	Panc vs. All: # peptides	1	0	0	0	3
6	# peptides unique to Panc	1	0	0	0	2
7	# false positives for Panc	0	n/a	n/a	n/a	1

All methods are able to successfully detect a dependence between peptide abundances in ovarian cancer samples versus normal samples (Table 3, line 1). This is likely because there are so many individual peptides that have different abundance distributions between ovarian and normal samples (Table 3, line 2). Nonetheless, Mgc identified more putative biomarkers than any of the other methods. While we have not checked all of them with subsequent experiments to identify potential false positives, we do know from previous experiments that three peptides in particular are effective biomarkers. All three peptides have p-value ≈ 0 for all methods including Mgc, that is, they are all correctly identified as significant. However, by ranking the peptides based on the actual test statistic of each peptide, Mgc is the method that ranks the three known biomarkers the lowest, suggesting that it is the least likely to falsely identify peptides.

We then investigated the pancreatic samples in an effort to identify biomarkers that are unique to pancreas. We first checked whether the methods could identify a difference using all the peptides. Indeed, three methods found a dependence at the 0.05 level, with Mgc obtaining the lowest p-value (Table 3, line 3). We then investigated how many individual peptides the methods identified; all of them found 6 to 11 peptides with a significant difference between pancreatic and normal samples (Table 3, line 4). Because we were interested in identifying peptides that were uniquely useful for pancreatic cancer, we then compared pancreatic samples to all others. At significance level 0.05, only Mgc, Hsic, and Hhg identified peptides that expressed different abundances in this more challenging case, and we list the top 4 peptides in Table 4 along with the corrected p-value for each peptide.

All three methods reveal the same unique protein for pancreas: neurogranin. Hsic identifies another peptide (tropomyosin alpha-3 chain isoform 4), and Hhg identifies a third peptide (fibrinogen-like protein 1 precursor). However, fibrinogen-like protein 1 precursor is not significant for p-value testing between

pancreatic and normal subjects. On the other hand, tropomyosin is a ubiquitously expressed protein, since normal tissues and other cancers will also express tropomyosin and leak it into blood, whereas neurogranin is exclusively expressed only in brain tissues. Moreover, there exists strong evidence of tropomyosin 3 upregulated in other cancers [94–97]. Therefore, it suggests that the other two peptides identified by HHG and HSIC are likely false positives.

In fact, neurogranin is always one of the top 4 candidates in all methods except MIC; the only difference is that the corrected p-values are not significant enough for other methods. Along with the classification result in Figure 4D showing that neurogranin alone has the best classification error, MGC discovers an ideal candidate for potential biomarker. Moreover, the fact that MGC, HHG and HSIC discover the dependency while others cannot implies a nonlinear relationship.

Table 4: For each of MGC, DCORR, MCORR, HHG, HSIC, MANTEL, PEARSON, and MIC, list the top 4 peptides identified for Panc vs All and the respective corrected p-value using Benjamini-Hochberg. Bold indicates a significant peptide at type 1 error level 0.05. The top candidates are very much alike except MIC. In particular, neurogranin is consistently among the top candidates for all methods, but is only significant while using MGC, HSIC, and HHG; there are two other significant proteins from HSIC and HHG, but they do not further improve the classification performance comparing to just using neurogranin. Note that the p-values from MANTEL and PEARSON are always 1 after Benjamini-Hochberg correction, so their respective top peptides are identified using raw p-values without correction.

Method	Top Four Identified Peptides			
MGC p-value	neurogranin 0.03	fibrinogen protein 1 0.33	tropomyosin alpha-3 0.49	ras suppressor protein 1 0.52
DCORR p-value	neurogranin 0.41	fibrinogen protein 1 0.60	kinase 6 0.60	twinfilin-2 0.93
MCORR p-value	neurogranin 0.45	fibrinogen protein 1 0.80	kinase 6 0.80	tropomyosin alpha-3 0.83
HSIC p-value	neurogranin 0.01	tropomyosin alpha-3 0.01	kinase 6 0.09	tripeptidyl-peptidase 2 0.09
HHG p-value	neurogranin 0.03	fibrinogen protein 1 0.03	tropomyosin alpha-3 0.03	platelet basic protein 0.11
MANTEL p-value	neurogranin 1	adenylyl cyclase 1	tropomyosin alpha-3 1	alpha-actinin-1 1
PEARSON p-value	neurogranin 1	adenylyl cyclase 1	tropomyosin alpha-3 1	alpha-actinin-1 1
MIC p-value	kinase B 0.15	S100-A9 0.15	ERF3A 0.15	thymidine 0.15

MGC Does Not Inflate False Positive Rates in Screening

In this final experiment, we empirically determine that MGC does not inflate false positive rates via a neuroimaging screening. To do so, we extend the work of Eklund et al. [48, 49], where a number of parametric methods are shown to largely inflate the false positives. Specifically, we applied MGC to test whether there is any dependency between brain voxel activities and random numbers. For each brain region, MGC attempts to test the following hypothesis: Is activity of a brain region independent of the time-varying stimuli? Any region that is selected as significant is a false positive by construction. By testing each brain region separately, MGC provides a distribution of false positive rates. If MGC is valid, the resulting distribution should be centered around the significance level, which is set at 0.05 for these

experiments. We considered 25 resting state fMRI experiments from the 1,000 Functional Connectomes Project consisting of a total of 1,583 subjects [98]. Figure 5 shows the false positive rates of Mgc for each dataset, which are centered around the critical level 0.05, as it should be. In contrast, many standard parametric methods for fMRI analysis, such as generalized linear models, can significantly increase the false positive rates, depending on the data and pre-processing details [48, 49]. Moreover, even the proposed solutions to those issues make linearity assumptions, thereby limiting detection to only a small subset of possible dependence functions.




Figure 5: We demonstrate that Mgc is a valid test that does not inflate the false positives in screening and variable selection. This figure shows the density estimate for the false positive rates of applying Mgc to select the “falsely significant” brain regions versus independent noise experiments; dots indicate the false positive rate of each experiment. The mean \pm standard deviation is 0.0538 ± 0.0394 .

Bibliography

- [1] J H Zhang, T D Chung, and K R Oldenburg. A simple statistical parameter for use in evaluation and validation of high throughput screening assays. *J. Biomol. Screen.*, 4(2):67–73, 1999.
- [2] Jeffrey William Prescott. Quantitative imaging biomarkers: the application of advanced image processing and analysis to clinical and preclinical decision making. *J. Digit. Imaging*, 26(1):97–108, February 2013.
- [3] T. Hastie, R. Tibshirani, and J. H. Friedman. *Elements of Statistical Learning*. Springer, New York, 2001.
- [4] Judea Pearl. *Causality: Models, Reasoning, and Inference*. Cambridge University Press, first edition edition, 13 March 2000.
- [5] K. Pearson. Notes on regression and inheritance in the case of two parents. *Proceedings of the Royal Society of London*, 58:240–242, 1895.
- [6] M. Reimherr and D. Nicolae. On quantifying dependence: A framework for developing interpretable measures. *Statistical Science*, 28(1):116–130, 2013.
- [7] J. Josse and S. Holmes. Measures of dependence between random vectors and tests of independence. *arXiv*, 2013. URL <http://arxiv.org/abs/1307.7383>.

- [8] W. Hoeffding. A non-parametric test of independence. *Annals of Mathematical Statistics*, 19(4): 546–557, 1948.
- [9] A. Renyi. On measures of dependence. *Acta Mathematica Academiae Scientiarum Hungarica*, 10(3):441–451, 1959.
- [10] J. Friedman and L. Rafsky. Graph-theoretic measures of multivariate association and prediction. *Annals of Statistics*, 11(2):377–391, 1983.
- [11] M. Schilling. Multivariate two-sample tests based on nearest neighbors. *Journal of the American Statistical Association*, 81(395):799–806, 1986.
- [12] G. Szekely, M. Rizzo, and N. Bakirov. Measuring and testing independence by correlation of distances. *Annals of Statistics*, 35(6):2769–2794, 2007.
- [13] G. Szekely and M. Rizzo. Brownian distance covariance. *Annals of Applied Statistics*, 3(4):1233–1303, 2009.
- [14] G. Szekely and M. Rizzo. The distance correlation t-test of independence in high dimension. *Journal of Multivariate Analysis*, 117:193–213, 2013.
- [15] R. Lyons. Distance covariance in metric spaces. *Annals of Probability*, 41(5):3284–3305, 2013.
- [16] A. Gretton and L. Györfi. Consistent nonparametric tests of independence. *Journal of Machine Learning Research*, 11:1391–1423, 2010.
- [17] Krikamol Muandet, Kenji Fukumizu, Bharath Sriperumbudur, and Bernhard Schölkopf. Kernel mean embedding of distributions: A review and beyond. *Foundations and Trends® in Machine Learning*, 10(1-2):1–141, June 2017.
- [18] D. Sejdinovic, B. Sriperumbudur, A. Gretton, and K. Fukumizu. Equivalence of distance-based and rkhs-based statistics in hypothesis testing. *Annals of Statistics*, 41(5):2263–2291, 2013.
- [19] Cencheng Shen and Joshua T Vogelstein. The exact equivalence of distance and kernel methods for hypothesis testing. June 2018.
- [20] R. Heller, Y. Heller, and M. Gorfine. A consistent multivariate test of association based on ranks of distances. *Biometrika*, 100(2):503–510, 2013.
- [21] Ronald Newbold Bracewell and Ronald N Bracewell. *The Fourier transform and its applications*, volume 31999. McGraw-Hill New York, 1986.
- [22] I Daubechies. *Ten lectures on wavelets*. SIAM, 1992.
- [23] Herbert Edelsbrunner and John L Harer. *Computational Topology: An Introduction*. American Mathematical Society, new ed. edition edition, December 2009.
- [24] Ronald R Coifman and Mauro Maggioni. Diffusion wavelets. *Appl. Comput. Harmon. Anal.*, 21(1): 53–94, July 2006.
- [25] William K. Allard, Guangliang Chen, and Mauro Maggioni. Multi-scale geometric methods for data sets II: Geometric Multi-Resolution Analysis. *Applied and Computational Harmonic Analysis*, 32(3):435–462, may 2012. ISSN 10635203. doi: 10.1016/j.acha.2011.08.001. URL <http://linkinghub.elsevier.com/retrieve/pii/S1063520311000868>.

- [26] Zhenyue Zhang, Jing Wang, and Hongyuan Zha. Adaptive manifold learning. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 34(2):253–265, 2012.
- [27] Arthur Gretton, Karsten M Borgwardt, Malte Rasch, Bernhard Schölkopf, and Alex J Smola. A kernel method for the two-sample-problem. In *Advances in neural information processing systems*, pages 513–520, 2006.
- [28] Elizaveta Levina and Peter J Bickel. Maximum likelihood estimation of intrinsic dimension. In *Advances in Neural Information Processing Systems*, 2004.
- [29] John A Lee and Michel Verleysen. *Nonlinear dimensionality reduction*. Springer Science & Business Media, 2007.
- [30] C. Shen, C. E. Priebe, and J. T. Vogelstein. From distance correlation to multiscale graph correlation. *Journal of the American Statistical Association*, 2018.
- [31] N. Mantel. The detection of disease clustering and a generalized regression approach. *Cancer Research*, 27(2):209–220, 1967.
- [32] H. Hotelling. Relations between two sets of variates. *Biometrika*, 28:321–377, 1936.
- [33] D. M. Witten, R. Tibshirani, and T. Hastie. A penalized matrix decomposition, with applications to sparse principal components and canonical correlation analysis. *Biostatistics*, 10(3):515–534, 2009.
- [34] D. M. Witten and R. Tibshirani. Penalized classification using Fisher’s linear discriminant. *Journal of the Royal Statistical Society, Series B*, 73(5):753–772, 2011.
- [35] C. Spearman. The Proof and Measurement of Association between Two Things. *The American Journal of Psychology*, 15(1):72, jan 1904. ISSN 00029556. doi: 10.2307/1412159. URL <http://www.jstor.org/stable/1412159?origin=crossref>.
- [36] M. G. Kendall. *Rank Correlation Methods*. London: Griffin, 1970.
- [37] D. Reshef, Y. Reshef, H. Finucane, S. Grossman, G. McVean, P. Turnbaugh, E. Lander, M. Mitzenmacher, and P. Sabeti. Detecting novel associations in large data sets. *Science*, 334(6062):1518–1524, 2011.
- [38] N. Simon and R. Tibshirani. Comment on detecting novel associations in large data sets. *arXiv*, 2012. URL <http://arxiv.org/abs/1401.7645>.
- [39] M. Gorfine, R. Heller, and Y. Heller. Comment on detecting novel associations in large data sets. *Technical Report*, 2012. URL <http://ie.technion.ac.il/~gorfinm/files/science6.pdf>.
- [40] Richard Berk, Lawrence Brown, Andreas Buja, Kai Zhang, Linda Zhao, et al. Valid post-selection inference. *The Annals of Statistics*, 41(2):802–837, 2013.
- [41] R. R. Costa, & McCrae. *Neo PI-R professional manual*, volume 396. 1992. ISBN 9997924452. doi: 10.1037/0003-066X.52.5.509.
- [42] Rex E. Jung, Judith M. Segall, H. Jeremy Bockholt, Ranee A. Flores, Shirley M. Smith, Robert S. Chavez, and Richard J. Haier. Neuroanatomy of creativity. *Human Brain Mapping*, 31(3):NA–NA, 2009. ISSN 10659471. doi: 10.1002/hbm.20874. URL <http://doi.wiley.com/10.1002/hbm.20874>.

- [43] Colin G DeYoung, Jacob B Hirsh, Matthew S Shane, Xenophon Papademetris, Nallakkandi Rajeevan, and Jeremy R Gray. Testing predictions from personality neuroscience: Brain structure and the big five. *Psychol. Sci.*, 21(6):820–828, June 2010.
- [44] Jiansong Xu and Marc N Potenza. White matter integrity and five-factor personality measures in healthy adults. *Neuroimage*, 59(1):800–807, January 2012.
- [45] Astrid Bjørnebekk, Anders M Fjell, Kristine B Walhovd, Håkon Grydeland, Svenn Torgersen, and Lars T Westlye. Neuronal correlates of the five factor model (FFM) of human personality: Multimodal imaging in a large healthy sample. *Neuroimage*, 65:194–208, January 2013.
- [46] Adriana Sampaio, José Miguel Soares, Joana Coutinho, Nuno Sousa, and Óscar F Gonçalves. The big five default brain: functional evidence. *Brain Struct. Funct.*, 219(6):1913–1922, November 2014.
- [47] J. Adelstein, Z. Shehzad, M. Mennes, C. DeYoung, X. Zuo, C. Kelly, D. Margulies, A. Bloomfield, J. Gray, F. Castellanos, and M. Milham. Personality is reflected in the brain’s intrinsic functional architecture. *PLoS ONE*, 6(11):e27633, 2011.
- [48] A. Eklund, M. Andersson, C. Josephson, M. Johansson, and H. Knutsson. Does parametric fmri analysis with spm yield valid results? an empirical study of 1484 rest datasets. *NeuroImage*, 61: 565–578, 2012.
- [49] Anders Eklund, Thomas Nichols, and Hans Knutsson. Cluster failure: Why fmri inferences for spatial extent have inflated false-positive rates. *Proceedings of the National Academy of Sciences*, 113(28):7900–7905, 2016.
- [50] Maria Frantzi, Akshay Bhat, and Agnieszka Latosinska. Clinical proteomic biomarkers: relevant issues on study design and technical considerations in biomarker development. *Clinical and Translational Medicine*, 3(7):PMC3994249, 2014.
- [51] Q. Wang, R. Chaerkady, J. Wu, H. J. Hwang, N. Papadopoulos, L. Kopelovich, A. Maitra, H. Matthaei, J. R. Eshleman, R. H. Hruban, K. W. Kinzler, A. Pandey, and B. Vogelstein. Mutant proteins as cancer-specific biomarkers. *Proceedings of the National Academy of Sciences of the United States of America*, (6):2444–9, 2011.
- [52] Kruttika Bhat, Fengfei Wang, Qingyong Ma, Qinyu Li, Sanku Mallik, Tze chen Hsieh, and Erxi Wu. Advances in biomarker research for pancreatic cancer. *Current Pharmaceutical Design*, 18(17): 2439–2451, 2012.
- [53] Q. Wang, M. Zhang, T. Tomita, J. T. Vogelstein, S. Zhou, N. Papadopoulos, K. W. Kinzler, and B. Vogelstein. A selected reaction monitoring approach for validating candidate biomarkers. *PNAS*, 2017.
- [54] Jun Yang, Frederick K. Korley, Min Dai, and Allen D. Everetta. Serum neurogranin measurement as a biomarker of acute traumatic brain injury. *Clinical Biochemistry*, 48:843–848, 2015.
- [55] E. Willemse, A. De Vos, E. Herries, U. Andreasson U, S. Engelborghs S, W. van der Flier, P. Scheltens, D. Crimmins, J. Ladenson, E. Vanmechelen, H. Zetterberg, A. Fagan, K. Blennow, M. Bjerke, and C. Teunissen. Neurogranin as cerebrospinal fluid biomarker for alzheimer disease: An assay comparison study. *Clinical Chemistry*, 64(6):927–937, 2018.

- [56] David Helfman, Patrick Flynn, Protiti Khan, and Ali Saeed. Tropomyosin as a regulator of cancer cell transformation. *Advances in Experimental Medicine and Biology*, 644:124–131, 2018.
- [57] C. Lam, C. Yip, T. Poon, C. Cheng, E. Ng, N. Wong, P. Cheung, P. Lai, I. Ng, S. Fan, and S. Cheung. Identification and characterization of tropomyosin 3 associated with granulin-epithelin precursor in human hepatocellular carcinoma. *PloS One*, 7(7):e40324, 2012.
- [58] Eric P Xing, Andrew Y Ng, Michael I Jordan, and Stuart Russell. Distance metric learning with application to clustering with side-information. *Advances in neural information processing systems*, 15:505–512, 2003.
- [59] Raja Giryes, Guillermo Sapiro, and Alex M Bronstein. Deep neural networks with random gaussian weights: A universal classification strategy. *CoRR*, abs/1504.08291, 2015.
- [60] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial nets. In *Advances in neural information processing systems*, pages 2672–2680, 2014.
- [61] Dougal J Sutherland, Hsiao-Yu Tung, Heiko Strathmann, Soumyajit De, Aaditya Ramdas, Alex Smola, and Arthur Gretton. Generative models and model criticism via optimized maximum mean discrepancy. In *International Conference on Learning Representations*, November 2016.
- [62] Gabor J. Szekely and Maria L. Rizzo. A new test for multivariate normality. *Journal of Multivariate Analysis*, 93(1):58–80, 2005. ISSN 0047259X. doi: 10.1016/j.jmva.2003.12.002.
- [63] Maria L. Rizzo and Gabor J. Szekely. DISCO analysis: A nonparametric extension of analysis of variance. *Annals of Applied Statistics*, 4(2):1034–1055, 2010.
- [64] Gábor J. Székely and Maria L. Rizzo. Partial distance correlation with methods for dissimilarities. *The Annals of Statistics*, 42(6):2382–2412, dec 2014. ISSN 0090-5364. doi: 10.1214/14-AOS1255. URL <http://projecteuclid.org/euclid.aos/1413810731>.
- [65] Xueqin Wang, Wenliang Pan, Wenhao Hu, Yuan Tian, and Heping Zhang. Conditional Distance Correlation. *Journal of the American Statistical Association*, 110(512):1726–1734, oct 2015. ISSN 0162-1459. doi: 10.1080/01621459.2014.993081. URL <http://www.tandfonline.com/doi/full/10.1080/01621459.2014.993081>.
- [66] R. Li, W. Zhong, and L. Zhu. Feature screening via distance correlation learning. *Journal of American Statistical Association*, 107:1129–1139, 2012.
- [67] Wei Zhong and Liping Zhu. An iterative approach to distance correlation-based sure independence screening. *Journal of Statistical Computation and Simulation*, 85(11):2331–2345, jul 2015. ISSN 0094-9655. doi: 10.1080/00949655.2014.928820. URL <http://www.tandfonline.com/doi/full/10.1080/00949655.2014.928820>.
- [68] Gabor J. Szekely and Maria L. Rizzo. Testing for Equal Distributions in High Dimension. *InterStat*, 10(11), 2004. doi: 10.1.1.226.377.
- [69] 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.

- [70] M. Rizzo and G. Székely. Energy distance. *Wiley Interdisciplinary Reviews: Computational Statistics*, 8(1):27–38, 2016.
- [71] J. R. Kettenring. Canonical analysis of several sets of variables. *Biometrika*, 58(3):433–451, 1971.
- [72] A. Tenenhaus and M. Tenenhaus. Regularized generalized canonical correlation analysis. *Psychometrika*, 76(2):257–284, 2011.
- [73] C. Shen, M. Sun, M. Tang, and C. E. Priebe. Generalized canonical correlation analysis for classification. *Journal of Multivariate Analysis*, 130:310–322, 2014.
- [74] J. B. Tenenbaum, V. de Silva, and J. C. Langford. A global geometric framework for nonlinear dimension reduction. *Science*, 290:2319–2323, 2000.
- [75] L. K. Saul and S. T. Roweis. Nonlinear dimensionality reduction by locally linear embedding. *Science*, 290:2323–2326, 2000.
- [76] X. Huo and G. Székely. Fast computing for distance covariance. *Technometrics*, 58(4):435–447, 2016.
- [77] C. Huang and X. Huo. A statistically and numerically efficient independence test based on random projections and distance covariance. *arXiv*, 2017.
- [78] Qinyi Zhang, Sarah Filippi, Arthur Gretton, and Dino Sejdinovic. Large-scale kernel methods for independence testing. *Statistics and Computing*, pages 1–18, 2017.
- [79] Arin Chaudhuri and Wenhao Hu. A fast algorithm for computing distance correlation. <https://arxiv.org/abs/1810.11332>, 2018.
- [80] Da Zheng, Disa Mhembere, Randal Burns, Joshua T Vogelstein, Carey E Priebe, and Alexander S Szalay. **FlashGraph: Processing Billion-Node Graphs on an Array of Commodity SSDs**. In *USENIX Conference on File and Storage Technologies*, 2015.
- [81] Da Zheng, Disa Mhembere, Joshua T. Vogelstein, Carey E. Priebe, and Randal Burns. **FlashMatrix: Parallel, Scalable Data Analysis with Generalized Matrix Operations using Commodity SSDs**. *arXiv*, 1604.06414, 2016. URL <http://arxiv.org/abs/1604.06414v1>.
- [82] Da Zheng, Randal Burns, Joshua T. Vogelstein, Carey E. Priebe, and Alexander S. Szalay. **An SSD-based eigensolver for spectral analysis on billion-node graphs**. *arXiv:1602.01421*, 2016.
- [83] Da Zheng, Disa Mhembere, Vince Lyzinski, Joshua T. Vogelstein, Carey E. Priebe, and Randal Burns. Semi-External Memory Sparse Matrix Multiplication on Billion-node Graphs in a Multicore Architecture. *arXiv*, feb 2016. URL <http://arxiv.org/abs/1602.02864>.
- [84] P. Good. *Permutation, Parametric, and Bootstrap Tests of Hypotheses*. Springer, 2005.
- [85] Gilles Guillot and François Rousset. Dismantling the mantel tests. *Methods in Ecology and Evolution*, 4(4):336–344, 2013.
- [86] G. Székely and M. Rizzo. Partial distance correlation with methods for dissimilarities. *Annals of Statistics*, 42(6):2382–2412, 2014.

- [87] Cameron Craddock, Sharad Sikka, Brian Cheung, Ranjeet Khanuja, Satrajit S Ghosh, Chaogan Yan, Qingyang Li, Daniel Lurie, Joshua Vogelstein, Randal Burns, Stanley Colcombe, Maarten Mennes, Clare Kelly, Adriana Di Martino, Francisco Xavier Castellanos, and Michael Milham. Towards Automated Analysis of Connectomes: The Configurable Pipeline for the Analysis of Connectomes (C-PAC). *Frontiers in Neuroinformatics*, 42, 2015. ISSN 1662-5196. doi: 10.3389/conf.fninf.2013.09.00042. URL <http://www.frontiersin.org/neuroinformatics/10.3389/conf.fninf.2013.09.00042/full>.
- [88] Danai Koutra, N Shah, Joshua T. Vogelstein, B J Gallagher, and Christos Faloutsos. DeltaCon: A Principled Massive-Graph Similarity Function. *ACM Transactions on Knowledge Discovery from Data*, may 2015.
- [89] William Gray Roncal, Zachary H. Koterba, Disa Mhembere, Dean M Kleissas, Joshua T. Vogelstein, Randal Burns, Anita R. Bowles, Dimitrios K. Donavos, Sephira Ryman, Rex E. Jung, Lei Wu, Vince D. Calhoun, and R. Jacob Vogelstein. MIGRAINE: MRI Graph Reliability Analysis and Inference for Connectomics. *Global Conference on Signal and Information Processing*, 2013. URL <http://arxiv.org/abs/1312.4875>.
- [90] Daniel L. Sussman, Minh Tang, Donniell E. Fishkind, and Carey E. Priebe. A consistent dot product embedding for stochastic blockmodel graphs. *Journal of the American Statistical Association*, 107 (499):1119–1128, 2013. doi: 10.1080/01621459.2012.699795. URL <http://arxiv.org/abs/1108.2228>.
- [91] C. Shen, J. T. Vogelstein, and C. Priebe. Manifold matching using shortest-path distance and joint neighborhood selection. *arXiv*, 2016. URL <http://arxiv.org/abs/1412.4098>.
- [92] Minh Tang, Avanti Athreya, Daniel L. Sussman, Vince Lyzinski, Youngser Park, and Carey E. Priebe. A semiparametric two-sample hypothesis testing problem for random graphs. *Journal of Computational & Graphical Statistics*, 2016. URL <http://www.tandfonline.com/doi/abs/10.1080/10618600.2016.1193505>.
- [93] Yoav Benjamini and Yosef Hochberg. Controlling the false discovery rate: a practical and powerful approach to multiple testing. *Journal of the royal statistical society. Series B (Methodological)*, pages 289–300, 1995.
- [94] S. Karsani, N. Saihen, R. Zain, S. Cheong, and M. Rahman. Comparative proteomics analysis of oral cancer cell lines: identification of cancer associated proteins. *Proteome Science*, (1):3, 2014.
- [95] Y. Sun, Z. Xia, Z. Shang, K. Sun, X. Niu, L. Qian, L. Fan, C. Cao, and H. Xiao. Facile preparation of salivary extracellular vesicles for cancer proteomics. *Scientific Reports*, page 24669, 2016.
- [96] H. Lee, C. Lim, Y. Cheong, M. Singh, and L Gam. Comparison of protein expression profiles of different stages of lymph nodes metastasis in breast cancer. *International Journal of Biological Sciences*, (3):353–362, 2012.
- [97] C. Y. Lam, C. W. Yip, T. C. Poon, C. K. Cheng, E. W. Ng, N. C. Wong, P. F. Cheung, P. B. Lai, I. O. Ng, S. T. Fan, and S. T. Cheung. Identification and characterization of tropomyosin 3 associated with granulin-epithelin precursor in human hepatocellular carcinoma. *PLoS ONE*, (7):e40324, 2012.

- [98] Bharat B Biswal, Maarten Mennes, Xi-Nian Zuo, Suril Gohel, Clare Kelly, Steve M Smith, Christian F Beckmann, Jonathan S Adelstein, Randy L Buckner, Stan Colcombe, et al. Toward discovery science of human brain function. *Proceedings of the National Academy of Sciences*, 107(10): 4734–4739, 2010.

Acknowledgment

This work was partially supported by the Child Mind Institute Endeavor Scientist Program, the National Science Foundation award DMS-1712947, the National Security Science and Engineering Faculty Fellowship (NSSEFF), the Johns Hopkins University Human Language Technology Center of Excellence (JHU HLT COE), the Defense Advanced Research Projects Agency's (DARPA) SIMPLEX program through SPAWAR contract N66001-15-C-4041, the XDATA program of DARPA administered through Air Force Research Laboratory contract FA8750-12-2-0303, DARPA Lifelong Learning Machines program through contract FA8650-18-2-7834, the Office of Naval Research contract N00014-12-1-0601, the Air Force Office of Scientific Research contract FA9550-14-1-0033. The authors thank Dr. Brett Mensh of Optimize Science for acting as our intellectual consigliere, Julia Kuhl for help with figures, and Dr. Ruth Heller, Dr. Bert Vogelstein, Dr. Don Geman, and Dr. Yakir Reshef for insightful suggestions. And we'd like to thank Satish Palaniappan, Sambit Panda, and Junhao (Bear) Xiong for porting the code to Python.

Supplementary Figures

Figure 2-Figure Supplements

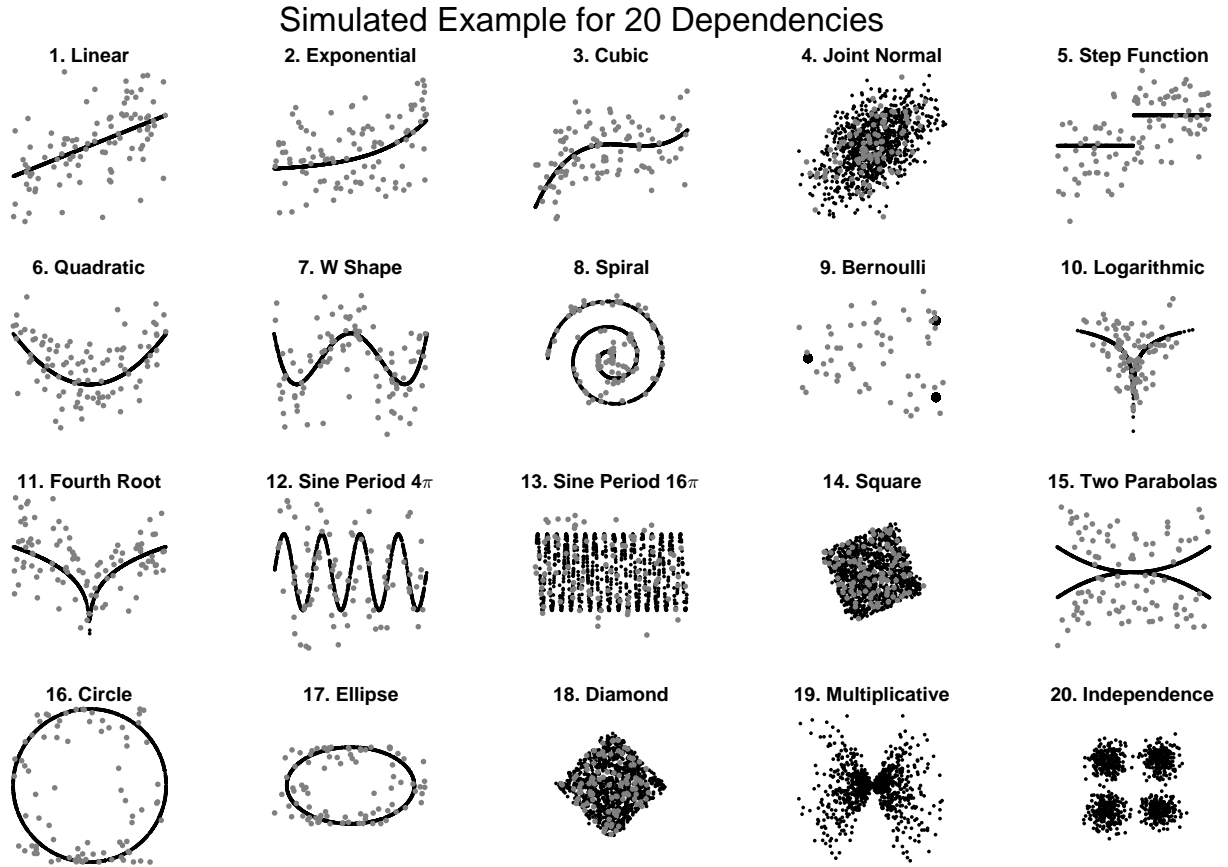


Figure 1: Visualization of the 20 dependencies at $p = q = 1$. For each, $n = 100$ points are sampled with noise ($\kappa = 1$) to show the actual sample data used for one-dimensional relationships (gray dots). For comparison purposes, $n = 1000$ points are sampled without noise ($\kappa = 0$) to highlight each underlying dependency (black dots). Note that only black points are plotted for type 19 and 20, as they do not have the noise parameter κ .

Testing Power for 20 Simulated 1-Dimensional Settings

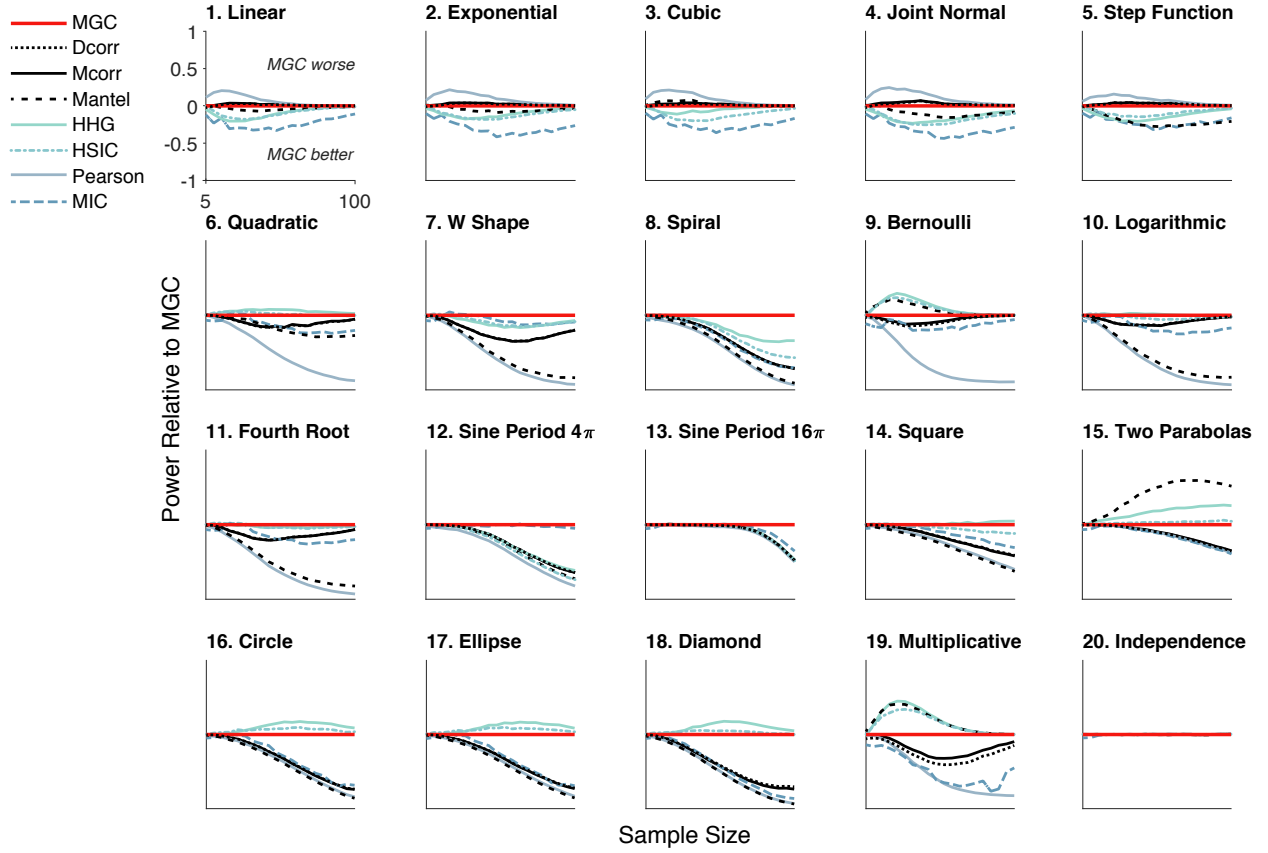


Figure 2: The same power plots as in Figure 2, except the 20 dependencies are one-dimensional with noise, and the x-axis shows sample size increasing from 5 to 100. Mgc empirically achieves similar or better power than the previous state-of-the-art approaches on most problems. Note that MIC is included in 1D case; RV and CCA both equal PEARSON in 1D; KENDALL and SPEARMAN are too similar to PEARSON in power and thus omitted in plotting.




Figure 3: The same set-ups as in Figure 2, comparing different Mgc implementations versus its global counterparts. The default Mgc builds upon Mcorr throughout the paper, and we further consider Mgc on MANTEL to illustrate the generalization. The magenta line shows the power difference between Mcorr and Mgc, and the cyan line shows the power difference between MANTEL and the Mgc version of MANTEL. Indeed, Mgc is able to improve the global counterpart in testing power under nonlinear dependencies, and maintains similar power under linear and independent dependencies.




Figure 4: The same power plots as in Figure 3, except the 20 dependencies are one-dimensional with noise, and the x-axis shows sample size increasing from 5 to 100.

Figure 3-Figure Supplements

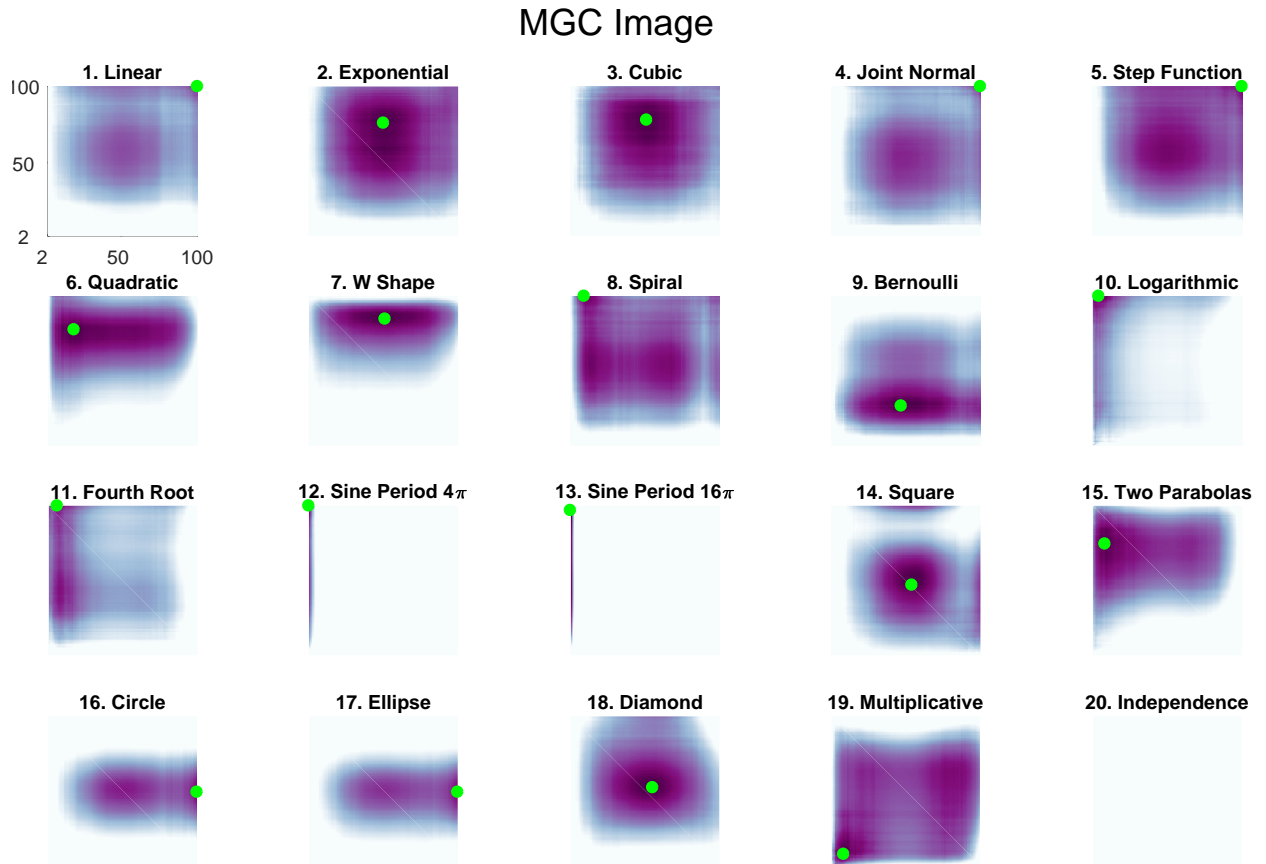


Figure 1: The Mgc-Map for the 20 panels for high-dimensional dependencies. For each simulation, the sample size is 100, and the dimension is selected as the dimension such that Mgc has a testing power above 0.5. It has similar behavior and interpretation as the 1-dimensional power maps in Figure 3, i.e., the linear relationships optimal scales are global, and similar dependencies share similar Mgc-Maps.

Running Time Report in Experiments

Here we list the actual running time of `MGC` versus other methods for testing on the real data, based on a modern desktop with a six core I7-6850K CPU and 32GB memory on Matlab 2017a on Windows 10. The first two experiments are timed based on 1000 permutations, while the screening experiment is timed without permutation, i.e., compute the test statistic only. `PEARSON` runs the fastest, trailed by `MIC` and then `DCORR`. `PEARSON` and `MIC` are only possible to run in the screening experiment, as the other two experiments are multivariate. The running time of `MGC` is a constant times (about 10) higher than that of `DCORR`, and `HHG` is implemented in a running time of $O(n^3)$ and thus significantly slower.

Table 5: The Actual Testing Time (in seconds) on Real Data.

Data	Personality	Creativity	Screening
<code>MGC</code>	2.5	7.5	1.9
<code>DCORR</code>	0.2	0.4	0.18
<code>HSIC</code>	0.5	1.7	0.23
<code>HHG</code>	6.3	53.4	12.3
<code>PEARSON</code>	NA	NA	0.03
<code>MIC</code>	NA	NA	0.1