

PAIRWISE NONLINEAR DEPENDENCE ANALYSIS OF GENOMIC DATA

BY SIQI XIANG^{1,a}, WAN ZHANG^{1,b}, SIYAO LIU^{2,e}, KATHERINE A. HOADLEY^{2,f},
CHARLES M. PEROU^{2,g}, KAI ZHANG^{1,c} AND J. S. MARRON^{1,d}

¹Department of Statistics and Operations Research, University of North Carolina at Chapel Hill, ^axiangsiqi.unc@gmail.com,
^bwan.zhang@unc.edu, ^czhangk@email.unc.edu, ^dmarron@unc.edu

²Lineberger Comprehensive Cancer Center, University of North Carolina at Chapel Hill, ^esiyao@email.unc.edu,
^fhoadley@med.unc.edu, ^gcperou@med.unc.edu

In The Cancer Genome Atlas (TCGA) data set, there are many interesting nonlinear dependencies between pairs of genes that reveal important relationships and subtypes of cancer. Such genomic data analysis requires a rapid, powerful, and interpretable detection process, especially in a high-dimensional environment. We study the nonlinear patterns among the expression of pairs of genes from TCGA using a powerful tool called binary expansion testing. We find many nonlinear patterns, some of which are driven by known cancer subtypes, some of which are novel.

1. Introduction. A leading cause of death in the world is cancer. A lot of cancer research is currently analyzing larger and larger data sets. In this paper we focus on The Cancer Genome Atlas (TCGA) (2012), which is a particularly important and comprehensive data set to study cancer biology and genomics. It is a publicly available genomics data resource that seeks to understand several types of cancer by collecting multiple diverse data over many people. In that set an important data type is gene expression and, more specifically, RNA-seq data. A biologically useful task of modern genomics data analysis is detecting the dependency patterns among gene expression. Conventional approaches to dependency, such as Pearson's, Spearman's rank, or Kendall's rank correlation coefficients target linear dependence. The focus of this paper is a much deeper investigation of nonlinear dependence in TCGA breast cancer data. The dependence of two such genes was an example shown in Zhang (2019). Here we carry the applied analysis much further by a detailed study of all pairs of genes. Furthermore, we take the analysis even deeper by studying gene dependence within subtypes as well. We show nonlinear dependence plays a much stronger role than previously imagined by finding 167,173 interesting nonlinear dependent pairs in the full data set. As seen in Table 4 below, only 36.3% of these significant pairs were discovered by using the classical Hoeffding's D statistic.

An example of the expression of two genes with strong and important nonlinear dependence that is not discoverable by linear methods is shown in the left panel of Figure 1. That is a scatter plot of expression for the genes BCL11A and F2RL2. For this pair of genes, the Pearson correlation coefficient is -0.0069 , Spearman's ρ is 0.088 , and Kendall's τ is 0.085 . These correlation coefficients are all very close to zero, suggesting no linear correlation between these two genes, as is also visually apparent. However, there is clear nonlinear dependence. This dependence is explained from a biological viewpoint by labeling with commonly used breast cancer subtypes, which were originally discovered by clustering some carefully selected genes in Perou et al. (2000). As shown by the colors and symbols in Figure 1, the decreasing part on the right (suggesting negative correlation) is mainly caused by the basal (Δ) subtype observations. For gene network considerations, the increasing part on the left,

Received June 2022; revised November 2022.

Key words and phrases. Binary expansion, genomic data, nonlinear dependence, nonparametric dependence testing.




FIG. 1. *Left: The scatter plot comparing expression of two genes in TCGA breast cancer data in the normalized log count scale. Right: The scatter plot of the same two genes using the copula transformation with the nonlinear dependence pattern from BET. Strong statistical significance is indicated by the BET Z-statistic of 12.56. These pairwise genes exhibit an interesting nonlinear dependence pattern, which is explained by showing the breast cancer subtypes.*

driven by the Luminal A (+) and Luminal B (\square) subtypes, suggest a positive correlation. This important biological refinement of network analysis is unavailable from classical gene network approaches based on conventional correlation measures. This nonlinear dependence pattern has been discovered by binary expansion testing (BET) proposed by Zhang (2019). The right panel of Figure 1 is the corresponding BET diagnostic plot explained in Section 2.1.

The data set studied in this paper consists of gene expression features of TCGA Lobular Freeze breast cancer data from Ciriello et al. (2015), containing 16,615 genes. Note that the total number of pairwise comparisons of genes is $\binom{16,615}{2} = 138,020,805$. Human visualization of all of these scatter plots is intractable (Sun and Zhao (2014)). An interesting early approach to this is Tukey's scagnostics (Wilkinson, Anand and Grossman (2005)). The large number of pairs motivates a computationally efficient method for investigating pairwise dependence. Gene expression studies have revealed a large number of linear dependencies between genes. In particular, in our BET analysis we discover 10,110,787 pairs of (statistically significant using Bonferroni multiple comparison adjustment) dependencies between genes. A large number of these are well understood. This paper takes genomics in a new direction by investigating nonlinear dependencies of the type shown in Figure 1.

There are several current approaches to studying nonlinear dependence. An early measure was Hoeffding's D (Hoeffding (1948)), which is not particularly powerful in the direction of nonlinear dependence (Zhang (2019)). More recently, Székely et al. have proposed the more powerful (in the direction of nonlinear dependence) method of distance correlation (Székely and Rizzo (2013), Székely, Rizzo and Bakirov (2007)). Another more powerful approach is the k -nearest neighbor mutual information (KNN-MI) algorithm (Kinney and Atwal (2014), Kraskov, Stögbauer and Grassberger (2004)), which focuses on mixtures of Gaussian distributions. While these methods are beneficial for discovering nonlinear dependence, they are less suitable for extensive genomic studies for three reasons. First, they are not efficient for large-scale computation problems such as TCGA data (tens of thousands of genes and hundreds of samples). Second, they still face some power loss in the direction of nonlinear dependence, as noted using simulation studies in Section 6 of Zhang (2019). Third, there is less immediate interpretation of the type that is available from BET.

To demonstrate the relatively slow execution time needed for the three methods above, we compare their calculation speed with BET. The running time for testing all pairs of a

TABLE 1

The running time of the pairwise comparison of 100 genes. The more powerful nonlinear detection methods are orders of magnitude slower than BET

Algorithms	BET	Hoeffding's D	Distance correlation	KNN mutual information
Times	8.96 secs	24.05 secs	17.51 mins	4.91 hours

randomly selected set of 100 genes is shown in Table 1. We use the default setting for each algorithm.

Table 1 shows that BET is around three times faster than Hoeffding's D, which has less power in the direction of nonlinear dependence, as seen in Section 3.2. It also shows that BET provides computational speed that is several orders of magnitude faster than either distance correlation or KNN mutual information in the context of dependence testing of high-dimensional data such as TCGA. Furthermore, nonlinear dependence can arise in many forms. As mentioned above, another advantage of BET is that it gives additional information on the form of nonlinear dependence, as illustrated in Figure 2.




FIG. 2. The first nine low-resolution dyadic binary interaction designs (BIDs) used by BET and one example pair of expression data shown for every BID, with S indicating the difference of counts in white and shade regions. Each BID is aimed at detecting a particular dependence relationship between two $[0, 1]$ uniform random variables. Size of S indicates the strength of nonlinear dependence in each BID. A number of biologically relevant patterns are shown.

This paper is organized as follows. Section 2 describes the main idea of the binary expansion testing (BET) algorithm. Section 3 details the BET analysis of TCGA data set, revealing several interesting nonlinear patterns. Section 4 studies the validation of some surprising TCGA results using an independent genomics data set. Section 5 concludes the article.

2. Binary expansion testing for nonlinear dependency detection. BET is a recent and innovative approach for dependence testing that is powerful for detecting pairwise nonlinear dependence. Furthermore, it provides a computationally fast investigation process for large-scale data sets, as shown in Table 1. Finally, BET gives a clear interpretation for some specific nonlinear dependence patterns. These patterns are formally introduced in Section 2.1. The BET algorithm and inference are given in Section 2.2.

2.1. *BET dependence patterns.* The first step of BET is the copula transform of the bivariate distribution to the unit square $[0, 1]^2$ using a marginal probability integral transformation of each variable in the pair, as detailed in Section 2.2. The key idea of BET is to partition the unit square into different patterns that indicate interesting types of dependence in terms of counts, that is, densities of observations, in different regions. For fast computation these patterns are dyadic in nature. The first few of these are shown in Figure 2, where unions of shade blocks represent one region and white blocks are the alternative. Each partition pattern is called a *binary interaction design (BID)* by Zhang (2019). In our genomics data analysis, we consider the nine dependence BIDs shown in Figure 2. Each BID corresponds to one BET dependence pattern. If there is no dependence between U and V , the observations are randomly distributed in $[0, 1]^2$. Dependence patterns are reflected by significant differences between these shade and white region counts of points (density of observations). For each form of BID, the difference is called S (the *symmetry statistic*). These counts are tested against the null hypothesis of no difference between them. The value of S is given for each BID in Figure 2. Notice that S is positive if the white part is denser than the shade and negative otherwise. For example, the white region in the upper-left panel of Figure 2 (the label A_1B_1 of that BID will be explained later) contains more points and captures a monotone upward dependence, which corresponds to the large positive $S = 745$. Had the linear dependence in the data been downward, there would have been greater density in the shade region, and S would be negative. An example where the shade region is denser appears in the BID in the upper-right panel, where $S = -395$. Each dependence pattern in Figure 2 is illustrated using a pair of genes that strongly exhibit the corresponding BID, particularly the pair with the maximal absolute value of S .

Figure 2 shows that BET captures several different nonlinear relationships for strongly dependent pairs of genes. Specific genes in these pairs are listed in Table 2, where Gene 1 is on

TABLE 2
Gene names for each pair in Figure 2

Location	Gene 1 (U)	Gene 2 (V)
Top-left	C17orf81	C17orf61
Top-middle	RPL24	RPL9
Top-right	JAM3	ANKS6
Middle-left	RPL9	RPL32
Center	PRR15	CA12
Middle-right	CDH5	ZNF883
Bottom-left	ANKS6	DCN
Bottom-middle	CT62	FAM174A
Bottom-right	PROSC	ASH2L

the horizontal axis and Gene 2 is on the vertical axis. As noted above, the symmetry statistic S for each BID is far from 0, indicating that strong dependence exists. For example, the middle-left panel highlights a surprising bimodal pattern that we will discuss in detail in Section 3.6 and 4.2. The bottom-left panel indicates dependence roughly following a downward opening parabola, which is captured by a high density in the shade region with corresponding negative S value. Left and right opening parabolas are captured by the BID in the upper-right, where this pair of genes has a leftward opening parabola. In contrast, the bottom-right panel looks quite different. It is fairly close to linear dependence but clearly not bivariate normal. There are unusual concentrations in the upper right and the lower left.

The prevalence of these mixture patterns and some of the others frequently turns out to be a consequence of known breast cancer subtypes. In particular, basal is known to be very distinct. The dependencies highlighted by a number of different BIDs in Figure 2 are clearly explained by the separation between basal (Δ) and the other subtypes. The two mixture patterns, in the bottom-left and the top-right BIDs, demonstrate how the mixture of basal and other subtypes can drive such patterns. The great difference between the basal subtype and the rest also drives the dependency patterns in the center (A_2B_2) and the bottom-middle ($A_1A_2B_2$) BIDs. It is important to remember that several BIDs can respond (i.e., have a significant absolute value of S) to a given dependence pattern in the data. However, we made the computational choice of only keeping the largest absolute S value for each pair of genes. Hence, the top 200 most significant genes, shown in the graphs in Figures 6 and 8, may miss some overlapping pairs of genes. Potential future work of interest would be to study the impact of this choice.

From classical statistical viewpoints, such as those based on sparsity and correlation analysis of dependence, this data set has a perhaps surprising amount of nonlinear dependence, as seen from the numbers of significant gene pairs shown in Figure 4.

2.2. BET algorithm and inference. Next, we formally introduce the testing procedure and notations. For each pair of genes, consider the sample data as pairs of variables $(X_1, Y_1), \dots, (X_n, Y_n)$. We view these as realizations of two random variables X and Y . BET is a fully nonparametric method based on the copula transformation, which is computed from the marginal CDFs. In particular, let $U = F_X(X)$ and $V = F_Y(Y)$, which are uniform on $[0, 1]$, and preserve the relative relationship between X and Y . Because the CDFs $F_X(X)$ and $F_Y(Y)$ are often unknown in practice, BET approximates them using the empirical CDF. Thus, the i th observation in the empirical copula is (\hat{U}_i, \hat{V}_i) whose marginal distribution is uniformly distributed on the equally spaced support points $\frac{1}{n}, \dots, \frac{n-1}{n}, 1$ on $[0, 1]$.

A key motivation for BET is that each decimal fraction number in the interval $[0, 1]$ has a binary representation. The computation of binary white and shade dyadic subinterval patterns that underlie Figure 2 is motivated by the useful probabilistic binary expansions of the continuous uniform random variables U and V . These binary expansions (Kac (1959)) are $U = \sum_{k=1}^{\infty} A_k/2^k$ and $V = \sum_{k'=1}^{\infty} B_{k'}/2^{k'}$, where $A_k \stackrel{\text{iid}}{\sim} \text{Bernoulli}(1/2)$ and $B_{k'} \stackrel{\text{iid}}{\sim} \text{Bernoulli}(1/2)$. Similarly, each observation in the empirical copula \hat{U}_i and \hat{V}_i also has a binary expansion: $\hat{U}_i = \sum_{k=1}^{\infty} \hat{A}_{k,i}/2^k$, and $\hat{V}_i = \sum_{k'=1}^{\infty} \hat{B}_{k',i}/2^{k'}$. Note that the binary expansion of an observation \hat{U}_i is the binary representation of this number. Thus, \hat{A}_k and $\hat{B}_{k'}$ can be regarded as the 0–1 indicator functions containing the randomness of \hat{U}_i and \hat{V}_i , for example, $\hat{A}_1 = I(\hat{U}_i \in (1/2, 1])$ and $\hat{A}_k = I(\hat{U}_i \in \bigcup_{j=1}^{k-1} ((2j-1)/2^k, 2j/2^k])$.

The shade and white regions underlying the BET statistic are based on the truncation of these binary expansions at some finite depths d_1 and d_2 , respectively, $U_{d_1} = \sum_{k=1}^{d_1} A_k/2^k$ and $V_{d_2} = \sum_{k'=1}^{d_2} B_{k'}/2^{k'}$. The discrete variables U_{d_1} and V_{d_2} take on at most 2^{d_1} and 2^{d_2} values. Hence, there are $2^{d_1+d_2} - 1$ binary variables resulting from interactions between A_k and $B_{k'}$.

These variables are sufficient statistics to study interesting dependence (Zhang (2019)). In order to present these interaction variables in the form of products and reflect dependence between these products, we use the binary variables $\dot{A}_k = 2A_k - 1$ and $\dot{B}_{k'} = 2B_{k'} - 1$ to replace A_k and $B_{k'}$. Thus, the interaction events between A_k and $B_{k'}$ can be written as the products $\dot{A}_k \dot{B}_{k'}$. For example, the events $\{A_1 = 1, B_1 = 1\}$ and $\{A_1 = 0, B_1 = 0\}$ lead to the same interaction event $\{\dot{A}_1 \dot{B}_1 = 1\}$. Out of these $2^{d_1+d_2} - 1$ interactions, there are $(2^{d_1} - 1)(2^{d_2} - 1)$ variables of product form $\dot{A}_{k_1} \dots \dot{A}_{k_r} \dot{B}_{k'_1} \dots \dot{B}_{k'_t}$ for some $r, t > 0$. We call these variables *cross interactions*; each of which results from the product of at least one \dot{A}_k and one $\dot{B}_{k'}$ and reflects a BET partition (the BID mentioned earlier) in $[0, 1]^2$; that is, each product results in one type of white and shade regions partition. For example, in Figure 2 the BID $A_1 B_1$ represents the cross interaction variable $\dot{A}_1 \dot{B}_1$, where the observation i in the white region reflects the event $\{\dot{A}_{1,i} \dot{B}_{1,i} = 1\}$ and in the shade region reflects the event $\{\dot{A}_{1,i} \dot{B}_{1,i} = -1\}$. Thus, the depth parameters d_1 and d_2 decide the amount and type of BIDs considered in BET. The choice $d_1 = d_2 = 2$ is the right resolution to find the most dependence patterns of biological interest (as the nine BIDs shown in Figure 2). We do not consider larger choices of d for two reasons. First, since $d = 2$ cuts each interval into quarters and $d = 3$ cuts each interval into eighths, most of these patterns from larger depths, which can not be seen at depth $d = 2$, capture relationships that are not expected to give useful insights into the dependence inherent to gene expression. Second, it would entail a considerable computational cost: when given $d_1 = d_2 = 3$, we have a total of 49 BIDs, including these nine patterns, meaning the computational cost is raised by a factor of more than five. Zhang (2019) gives more discussion about depth selection. Let \mathbf{a} and \mathbf{b} denote vectors of length d_1 and d_2 with 1's at $k_1 \dots k_r$ and $k'_1 \dots k'_t$, respectively, and 0's otherwise; thus, we can denote the cross interaction $\dot{A}_{k_1} \dots \dot{A}_{k_r} \dot{B}_{k'_1} \dots \dot{B}_{k'_t}$ as $\dot{A}_a \dot{B}_b$. For simplicity of notation, the labels $A_a B_b$ in most figures represent the cross interactions $\dot{A}_a \dot{B}_b$.

Now, we define the symmetry statistic S for a given cross interaction $\dot{A}_a \dot{B}_b$ as the difference of counts in white and shade regions in the corresponding BID. Since the value of the cross interaction of the observation is 1 in white regions and -1 in shade regions, we can calculate S as the sum of the observed binary interaction variables $S = \sum_{i=1}^n \dot{A}_{a,i} \dot{B}_{b,i}$. The U_{d_1} and V_{d_2} are strongly dependent when the absolute value $|S|$ is far from zero for at least one BID, according to the following fundamental observation of Zhang (2019): If U_{d_1} and V_{d_2} are independent, the symmetry statistic S satisfies $(S + n)/2 \sim \text{Binomial}(n, 1/2)$, for $a \neq 0$ and $b \neq 0$. If the empirical copula transformation is used, we can use $\hat{S} = \sum_{i=1}^n \hat{A}_{a,i} \hat{B}_{b,i}$ as the symmetry statistic and $(\hat{S} + n)/4 \sim \text{Hypergeometric}(n, n/2, n/2)$, for $\mathbf{a} \neq 0$ and $\mathbf{b} \neq 0$.

The issue of multiple comparisons across BIDs is handled by the Max BET procedure of Zhang (2019) at depths d_1 and d_2 . This is described as follows for a given pair of variables. First, we compute all symmetry statistics S with cross interactions for the given d_1 and d_2 . Then we look for the symmetry statistic with the strongest asymmetry and record its p -value and z -statistic $|S|/\sqrt{n}$. Finally, for this pair of variables we use a Bonferroni adjustment across the cross interactions (BIDs) to obtain the corresponding familywise error rate p -value for this maximum $|S|$. The dependence relationship is represented by the most significant BID.

Notice that Figure 2 reveals a reflection property among the BIDs. For example, the BIDs $A_1 B_1 B_2$ (top-right) and $A_1 A_2 B_1$ (bottom-left) represent the same relationship up to a switch of the axis positions (i.e., reversal of the roles of the two genes) for the two random variables U and V . There are three pairs of such off-diagonal reflected patterns. Identifying such pairs results in six BID patterns: five nonlinear and one linearity. This reflection property will be further discussed in Section 3.3.




FIG. 3. *Left: The scatter plot comparing expression of two genes in the TCGA breast cancer data in the normalized log count scale with zero value imputation by the median; Middle: The scatter plot of the same two genes in the normalized log count scale with the median values reset to the minimum and jittering applied to nonunique values; Right: The scatter plot of the same two genes in the copula distribution scale.*

3. Results from TCGA. In this section we first expand on the data preprocessing of TCGA data set in Section 3.1. Then we summarize the analysis results in Section 3.2. Finally, we discuss some specific nonlinear dependency patterns in the last few subsections.

3.1. Data preprocessing. The RNA-seq gene expression features of TCGA Lobular Freeze breast cancer data set from Ciriello et al. (2015) contain 16,615 genes of 817 primary tumor samples, including five subtypes (proportion in the sample): Basal-like (16.6%), HER2 (8.0%), Luminal A (50.8%), Luminal B (21.5%), and Normal-like (3.1%). Intrinsic breast cancer subtyping was done using the PAM50 classifier (Parker et al. (2009)).

This gene expression data was preprocessed as described in Ciriello et al. (2015), which included normalization and logarithms. In particular, during this genomics data preprocessing each sample was normalized to a fixed upper quartile and then log 2 transformed. Genes with more than 20% zero counts were excluded. Other genes with zero counts had their zeros recorded as missing. A questionable choice made in that preprocessing was to replace these missing values by the median for that gene. However, such data was the beginning of our analysis. The poor consequences of this approach are illustrated in the left panel of Figure 3. In particular, it shows the raw data of a pair of example genes in our TCGA data set, which has many points piled up at the median representing missing values (zero counts). This is an inappropriate way of handling the zeros, because all these data are based on counts, so a zero count represents a small level of gene expression. As an alternative we first considered moving the median values to be the same as the smallest value. In the case of no zero counts, this is inappropriate because the median corresponds to a nonzero count. A simple fix to this is to take the first nonzero count and leave it at the median. Thus, we address this by setting all but the first of the median values to the minimum value. Some experimentation revealed that this has a minimal impact. The result of this process is shown in the middle panel of Figure 3. This causes many ties for the smallest value. Since BET is based on a copula transform that essentially assumes continuous variables, a large number of ties will cause severe noncontinuity and strongly impact the BET inference. Therefore, these points are spread out in the interval of the minimum and the second unique minimum, using a jitter approach. Specifically, to preserve the ranks, a small random value (uniformly distributed between 0 and the difference between the second unique minimum and minimum) is added to the nonunique minimum observations. Jittering has no impact on the BET significance because the jitter points are within the first 20% of the data (recall that genes with more than 20% missing were excluded). Finally, we apply the empirical copula transformation; see the right panel in Figure 3.

TABLE 3
Four contexts and the corresponding sample sizes

Contexts	All five subtypes	LumA/LumB/Her2	Her2/LumB	LumA
Sample sizes	817	656	241	415

As noted in Parker et al. (2009), these subtypes play a critical role in many aspects of breast cancer. Our analysis in Section 3.3 confirms that the basal subtype tends to be quite different from the rest. Hence, we also investigate the context of LumA/LumB/Her2, which tends to be dominated by the most numerous subtype LumA. We further investigate the non-LumA group, meaning the union of Her2/LumB, and LumA separately. Hence, our analysis focuses on the four different contexts shown in Table 3, which are chosen to highlight important aspects. Note that the BET analysis focuses on detecting nonlinearity but not subtypes. However, there is a relationship between nonlinearity and subtypes. We set the depth parameters in the BET algorithm d_1, d_2 to be 2, and only focus on five nonlinear BIDs, as discussed in Section 2.2. To control for multiple comparisons, in each context we use the Bonferroni adjustment across genes to modify the BET output p -value of each pair, which has already been adjusted across the nine BIDs. Specifically, we use the total number of pairwise comparisons of genes in TCGA (138,020,805) to do the Bonferroni correction. Then we use the level 0.05.

The BET detection results of TCGA breast cancer data for these four different contexts and five nonlinear BIDs are summarized in Section 3.2. More detailed descriptions are in Section 3.3 through Section 3.6.

3.2. Summary of BET analysis. As we discussed at the end of Section 2.2, there are three pairs of reflected nonlinear BIDs in all nine BIDs. Identifying such pairs results in five nonlinear BIDs; see the five columns of Figure 4. From this point on, the *Parabolic* BID shown in the first column refers to the union of this $A_1A_2B_1$ BID and its reflection $A_1B_1B_2$. In our TCGA analysis, this Parabolic BID frequently finds a mixture data pattern. Similarly, the *W* BID in the second column refers to the union of this A_2B_1 BID and its reflection A_1B_2 , since the shape of this BID looks like the letter W. This W BID tends to find a particular bimodal pattern in our TCGA data. The rows of Figure 4 are the four contexts. The top gene pairs are shown for each. The number of significant pairs for each BID is displayed at the top of each panel. No pair is shown for the Her2/LumB context with the BID $A_1A_2B_2$ because there is no significant pair of genes.

Figure 4 shows the Parabolic BID in the first column, which contains the largest number of significant pairs for each context. In those cases some obviously show mixtures of different subtype distributions, such as the triangle Basal subtype in the lower right region of the top five subtypes panel and the asterisk Her2 cases in the lower right region of the Her2/LumB panel. The other relationships look strong but are not explained by subtypes, perhaps motivating additional genomic research. In particular, gene expression is indicative of many biological phenomena. Some are related to cancer subtypes, and some are not. A more detailed discussion of the gene pair in the upper left appears in Section 3.3 and of the gene pair in the lower left in Section 3.4.

Some surprising bimodal patterns in column 2 are captured by the W BID, which seem not to be driven by the breast cancer subtype information. In Section 3.6 we discuss the gene pair on the top row of column 2. To investigate the potential biological relevance of this bimodal dependence, an independent data set is used and discussed in Section 4.2.

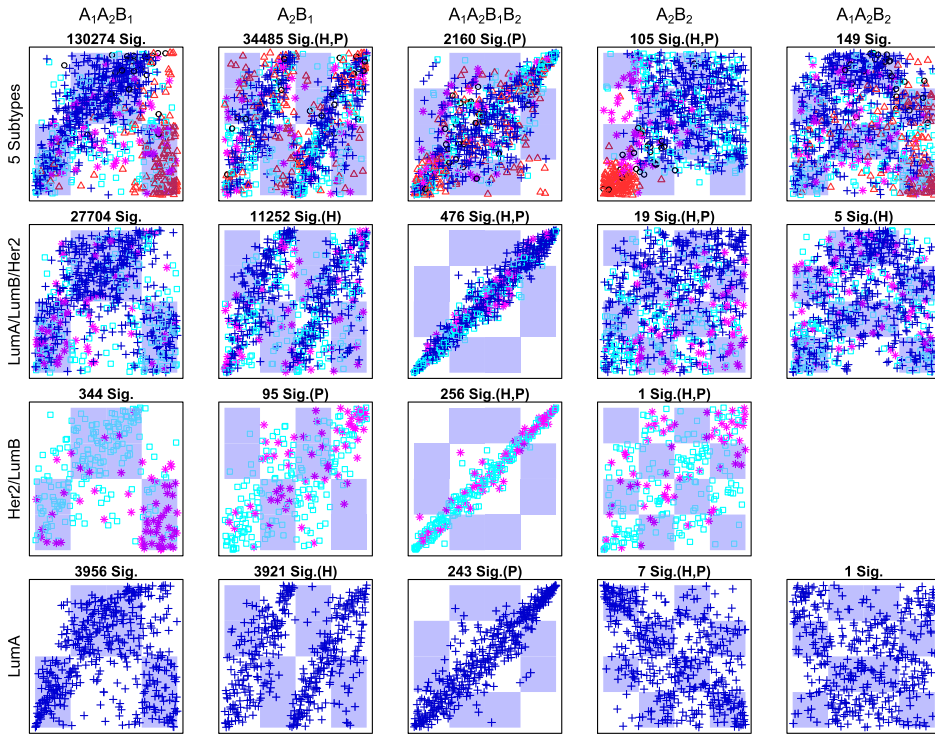


FIG. 4. *BET diagnostic plots where the four contexts are the rows and the four nonlinear BIDs are the columns. The most significant pair is shown for each. The number of significant pairs for each BID is shown at the top. The symbols “H” and “P” at the top show that this most significant pair is Hoeffding’s D and Pearson correlation, respectively, Bonferroni significant.*

In column 3 some approximately linear pairs of genes are detected by the BID $A_1 A_2 B_1 B_2$. In Section 3.5 we analyze the top pair in column 3 to discuss the connection between linearity and this nonlinear dependence pattern.

In column 4, as discussed in Section 2.1, we find an interesting biological separation in the five subtypes panel (top), which exists between the triangle basal subtype and other breast cancer subtypes. Moreover, in the LumA context (bottom), this checkerboard pattern suggests that, while many points are along the main diagonal, two small clusters lie off the main diagonal. This dependence pattern is not explained by the breast cancer subtypes and might be worth deeper biological investigations.

In column 5, the biological explanation of the most significant gene pairs, shown in each context, is not particularly clear. However, in Figure 2 the bottom-middle panel gives another example of this BID, which does contain an interesting biological pattern. We notice that the basal points are clustered in the upper-left corner, which are well separated from other subtypes. Furthermore, the Luminal B (square) subtype appears mostly at a cluster in the bottom center.

Figure 4 also allows the comparison of BET with Hoeffding’s D and Pearson correlation. In particular, the symbols “H” and “P” at the top of each panel indicate that the corresponding most significant gene pair is also Hoeffding’s D independence/Pearson correlation testing Bonferroni significant across all 1.38×10^8 possible pairs. From Figure 4, we find that many interesting nonlinear pairs of genes, discovered by BET, are not Bonferroni significant when using Hoeffding’s D or Pearson correlation, especially some biologically interesting parabolic patterns (Column 1). Hoeffding’s D and Pearson tend to discover approximately linear dependence patterns. Based on this observation, an interesting question is how many of these

TABLE 4

Comparison of BET, Hoeffding's D and Pearson Correlation over BET significant pairs of five BIDs in the five subtypes context

BIDs	$A_1 A_2 B_1$	$A_2 B_1$	$A_1 A_2 B_1 B_2$	$A_2 B_2$	$A_1 A_2 B_2$	Total
Number of BET significant pairs	130,274	34,485	2160	105	149	167,173
Number of Pearson correlation significant pairs	38,834	34,064	1662	105	98	74,763
Proportion of Pearson correlation significant pairs	29.8%	98.8%	76.9%	100.0%	65.8%	44.7%
Number of Hoeffding's D significant pairs	41,292	18,228	1071	56	58	60,705
Proportion of Hoeffding's D significant pairs	31.7%	52.9%	49.6%	53.3%	38.9%	36.3%

nonlinear significant dependence pairs found by BET can not be discovered by conventional methods. This is studied in Table 4. The pairs discovered by BET are assessed for significance by both classical linear Pearson correlation and the nonlinear Hoeffding's D. Recall that, according to Table 1, Hoeffding's D is about three times slower than BET. We compare these three methods over the significant pairs of genes of five nonlinear BIDs shown in Figure 4 in the five subtypes context. Specifically, for the Parabolic BID ($A_1 A_2 B_1$) in the five subtypes context, we apply Hoeffding's D over the 130,274 BET significant pairs of genes and record the count of Pearson correlation and Hoeffding's D significant results after the Bonferroni adjustment for p -values across all pairs. Then we calculate the proportion of Pearson correlation and Hoeffding's D significant pairs in the BET significant pairs in each BID. Table 4 summarizes the number of BET significant pairs (first row), the number (second row), and the proportion (third row) of Pearson correlation significant pairs as well as the number (fourth row) and the proportion (fifth row) of Hoeffding's D significant pairs for each BID.

Note that less than 30% of the pairs discovered by the Parabolic BID ($A_1 A_2 B_1$) were detected by Pearson correlation. For other patterns the performances of Pearson correlation are better. On the other hand, about the comparison with Hoeffding's D we notice that only 30%–50% BET significant pairs are detected by Hoeffding's D test in each of the five BIDs. Overall, 92,410 and 106,468 BET Significant Pairs are not discoverable by Pearson correlation and Hoeffding's D. This result indicates that BET is substantially more powerful against these particular types of nonlinear dependence.

3.3. Mixture pattern for five subtypes. In this section we discuss the pair of genes shown in the top panel of Figure 5: ANKS6 and JAM3, which have the most significant mixture dependence pattern. The dependence of these two genes is captured by the shade region of this BID. From the log-scale scatter plot (top-right) and BET diagnosis plot (top-left) in Figure 5, it is visually apparent that the basal subtype group is separated from the other subtypes. In particular, the nonbasal cases look like a classical bivariate Gaussian distribution with a *positive correlation* between ANKS6 and JAM3 (larger values of ANKS6 lead to more expression of JAM3). However, the basal cases behave very differently: larger ANKS6 goes along with smaller JAM3, indicating a *negative correlation*. Ignoring this important difference in direction of correlation can have a serious impact on gene network analysis. ANKS6 in this example is the main driver of the separation between basal and the others.

Here we revisit the reflection issue from the end of Section 2.2. As discussed above, the BIDs $A_1 A_2 B_1$ and $A_1 B_1 B_2$ are identical if we switch the genes on the x -axis and the y -axis. Since the pair of genes in this example (ANKS6 and JAM3) are detected by the $A_1 A_2 B_1$

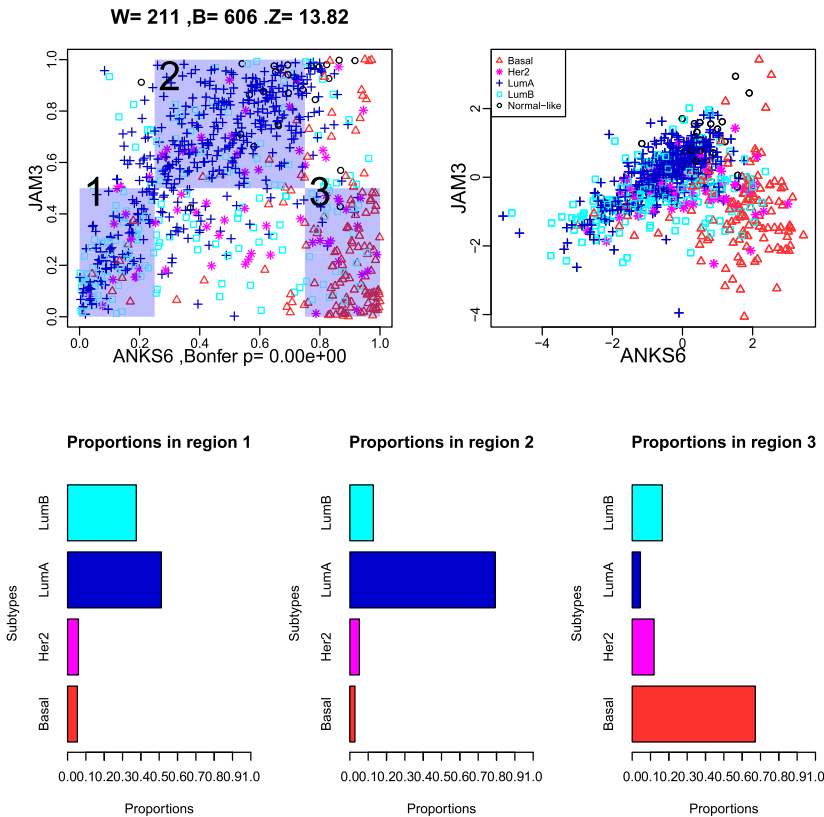


FIG. 5. Top-left: BET diagnosis plot for the pair of genes: ANKS6 and JAM3, which shows a strong surprising dependency with large z-statistic value and three shade rectangular regions. Top-right: The scatter plot of the same two genes in the normalized log count scale after preprocessing. Bottom: The proportions for each breast cancer subtype (except Normal-like) in each shade rectangular region.

BID, reversing the ordering of these two genes will give the reflected pattern ($A_1 B_1 B_2$). This example explains why only five nonlinear BIDs are considered here.

Further insight comes from splitting the shade region into three rectangular regions (see the numerical labels in the top-left panel) and calculating the respective proportions of the four breast cancer subtypes (ignoring the Normal-like) in each shade rectangle (bottom panels). These proportions reveal how the subtypes drive this relationship.

In the proportion bar plots, LumA and LumB subtypes have high proportions in Regions 1 and 2. Those two subgroups account for 88.7% and 92% in the two regions separately, where the basal subgroup only accounts for 5.4% and 2.8%. However, in Region 3 the proportion of basal reaches 67%, and the total proportion of LumA and LumB is 21%. This observation indicates that the positive correlation of the LumA and LumB domains are in Regions 1 and 2, and the differing basal correlation is in Region 3. The Pearson Chi-square test of independence is used to confirm this observation, for the counts of points in Table 5.

The p -value of this Chisq test is smaller than 2.2×10^{-16} (i.e., smaller than floating-point round-off error), and it shows a strong significance that these four subtypes are not homogeneously distributed in the shade regions. To more directly validate the separation between basal and the others, all others are combined into a single group, and the Chi-square test gives another small p -value less than round-off error. This result confirms the observation that this dependence pattern captured by the shade regions is very strongly significant and is influenced by the mixture of basal and other subtype distributions.

TABLE 5
Observed counts for four breast cancer subtypes in three shade regions

Subtypes	Region 1	Region 2	Region 3
Basal	9	8	90
LumA	86	229	6
LumB	63	37	22
Her2	10	15	16

The relationships between pairs of genes with respect to this same BID $A_1 B_1 B_2$ is shown by a network connection plot in Figure 6 (a similar network plot with respect to five nonlinear BIDs is given in Supplement A (Xiang et al. (2023))). Nodes in Figure 6 represent genes. This particular pattern of nonlinear BET dependence is highlighted by edges, which show Bonferroni statistical significance between genes. Furthermore, genes are ranked by their maximum BET z -scores. As noted in Zhang (2019), z -scores, reflecting the number of standard deviations above the mean, are more interpretable when p -values are extremely small. As shown in Figure 4, there are 130,274 significant pairs of genes for this BID and its reflection. To avoid a too cluttered network graphic, only the top 200 genes are shown in Figure 6. These 200 genes have 311 significant dependence edges for this Parabolic BID. Genes at the center of some visually important communities are labeled. Each community is a set of genes that shows this relationship with the center gene. Notice there are a number of gene communities representing different biological dependencies that are significant with respect to this BID. Figure 5 suggests that much of this nonlinear dependence may be due to the basal subtype, which is well known to be quite different from the others. However, there can be other causes of this pattern. For example, Figure 8 shows that the gene ZDHHC2 has such dependence, even when the basal subtype is left out of the analysis. This gene appears in Figure 6 as the point represented by the black triangle.

Good insights into any of these gene communities and their functions come from finding where they appear among published gene signatures, such as those shown in The Molecular Signatures Database (MSigDB), a collection of annotated gene sets for use with Gene Set Enrichment Analysis (Liberzon et al. (2011, 2015), Subramanian et al. (2005)). For example, we performed gene set enrichment analysis on the labeled communities. The gene set in the




FIG. 6. Network connection plot of 200 most significant genes (nodes) with 311 edges for the BID $A_1 B_1 B_2$. Each node represents one gene, and each edge represents a significant dependence between those genes. The large community illustrated there are many genes having significant mixture dependence pattern with the center gene, such as ANKS6. The point represented by the black triangle is the gene ZDHHC2 featured in Figure 8.




FIG. 7. *Left: BET diagnosis plot for a pair of genes which shows the mixture dependence within only the Luminal A subgroup. Right: The scatter plot of the same two genes in the normalized count scale. A point cloud with a strong positive correlation is on the left of the entire group, and the remaining cases form a more diffuse cluster on the bottom-right.*

largest (ANKS6) community is strongly associated with stromal or immune features. This is consistent with the previous finding that basal-like breast cancer has increased immune signature expression (Iglesia et al. (2016)). The gene analysis of the other label communities did not give such good biological interpretations. A full list of gene set enrichment analysis results for ANKS6 is in Supplement B (Xiang et al. (2023)).

3.4. Mixture pattern for only Luminal A subtype. While subtypes have played an important role in the diagnosis and treatment of breast cancer, the heterogeneity of the disease motivates deeper investigation within subtypes. Here we focus only on the Luminal A breast cancer subtype observations. Figure 7 shows an additional interesting mixture dependence pattern in both the BET diagnosis (left panel) and log-scale scatter (right panel) plots. The right panel contains a positively correlated Gaussian point cloud on the left. There is a more diffuse cluster toward the lower right. This seems to indicate a mixture behavior. In particular, the gene ZDHHC2 bifurcates the data into a cluster where it is strongly positively correlated with CELF2 and another cluster where large values of ZDHHC2 correspond to small values of CELF2. Hence, this pair of genes highlights potentially interesting subgroups, which merits a deeper investigation.

The connection plot in Figure 8 shows which genes have many significant pairs within the 200 most significant genes in the LumA only context with the Parabolic BID. There are 190 significant dependence edges. This can be less than 200 because there are many pairs that only connect with each other. ZDHHC2 and FGF10 are two central genes having large communities, which motivate a deeper investigation. Checking carefully, the individual plots reveal that, in all of these pairs, ZDHHC2 and FGF10 play the bifurcating role shown in Figure 7 in the dependence with each of these other genes. To further illustrate this bifurcation property, we show more examples of pairs about ZDHHC2 for this pattern in Supplement C (Xiang et al. (2023)).

To investigate the corresponding gene function, we use gene set enrichment analysis again to compute overlaps between these communities and gene sets from MSigDB. This ZDHHC2 community has overlaps with some gene sets related to breast cancer and, in particular, the luminal subtype, such as CHARAFE_BREAST_CANCER_LUMINAL_VS_MESENCHYMAL_DN and CHARAFE_BREAST_CANCER_LUMINAL_VS_BASAL_DN (Charafe-Jauffret et al. (2006)). A full list of overlapping results for ZDHHC2 is in Supplement D




FIG. 8. Connection plot of 200 most significant genes (nodes) with 190 edges for the Parabolic BID in the context of only Luminal A. ZDHHC2 and FGF10 are central genes in two large communities. The ZDHHC2 gene community has overlaps with some luminal gene sets in The Molecular Signatures Database (MSigDB).

(Xiang et al. (2023)). This confirms that ZDHHC2 and its community are an important player at luminal breast cancer and could motivate a deeper investigation into the role played by the ZDHHC2 community. On the other hand, a similar investigation of FGF10 doesn't show the connection with research to date on luminal breast cancer, again possibly motivating further biological work.

3.5. Connection between linear and nonlinear patterns. In column three of Figure 4, some approximately linear pairs of genes with a second-order structure are detected by the BID $A_1A_2B_1B_2$. The left panel of Figure 9 gives the BID diagnosis of the top row pair of genes in Column 3 (PROSC and ASH2L) with the largest symmetry statistic value and z-score in all nine BIDs. As discussed in Section 2.1, this BET diagnosis plot shows that this pair is not bivariate normal with concentrations in both ends of the diagonal. A deeper investigation of the structure of this pair of genes is from the scatter plot in the middle panel of Figure 9. It reflects an approximate linear relationship with strong skewness along the major axis. To analyze the connection between this pair and the linear dependence BID A_1B_1 , we show the linear BID in the right panel. The corresponding counts of white and shade regions for BID $A_1A_2B_1B_2$ in the left panel are 607 and 210 so that the symmetry statistic S and the z-score are 397 and 13.89; the corresponding counts for the linear BID in the right panel are 599 and 218 so that the symmetry statistic S and the z-score are rather close but slightly smaller values of 381 and 13.33. These numbers reflect the unusual pattern of greater




FIG. 9. Left: The most significant BID diagnosis plot for a pair of genes: PROSC and ASH2L. Middle: The scatter plot of the same two genes in the normalized count scale. Right: The corresponding linear BID A_1B_1 for the same two genes. The value of the BET statistic for the linear BID is slightly smaller than that for the most significant BID.




FIG. 10. *Left: BET diagnosis plot for a pair of genes that shows a bimodal dependence pattern. Right: The scatter plot of the same two genes in the normalized count scale. This nonlinear relationship seems to be driven by the bimodal distribution of gene RPL9.*

variation in the middle of the distribution with relatively less variation from the diagonal for the rest of it. This suggests a different type of mixture model, which may merit deeper investigation.

3.6. Bimodal pattern. The second column of Figure 4 shows a perhaps surprising bimodal dependence pattern that is shared by many pairs of genes.

To more deeply investigate this bimodal dependence, we take the top pair in the context of five subtypes as an example. Figure 10 shows the BET diagnosis (left panel) and log-scale scatter (right panel) plots of the most significant pair of this pattern in the context of five subtypes: RPL9 and RPL32. RPL9 in this pair separates the group into two positively correlated Gaussian clusters, suggesting this surprising bimodal dependence pattern perhaps is related to the gene RPL9. To investigate whether this is an important biological phenomenon or an artifact of some particular preprocessing steps in the Lobular Freeze TCGA data, we consider an additional completely separate data set in Section 4.2.

4. Biological reproducibility of TCGA results. An essential issue with exploratory data analyses, as shown in Section 3, is their reproducibility. To investigate this for the interesting results shown in Sections 3.3 and 3.6, we consider an independent genomic data set: the Sweden Cancerome Analysis Network-Breast (SCAN-B) (Brueffer et al. (2018)). This data set came from the NCBI Gene Expression Omnibus (GSE96058). The Data set was pre-processed, as described in (Saal et al. (2015)). We use a subset of the gene expression data set, which contains 2969 samples with full clinical data and 30,865 genes. There are 15,197 genes existing in both the SCAN-B and TCGA data sets. We only consider these common genes during this validation process. There was no further processing step in the SCAN-B set for our reproducibility analysis. First, in Section 4.1 we study the reproducibility of the mixture patterns in the contexts of the five subtypes, as shown in Section 3.3. Then in Section 4.2, we find that the bimodal distribution of the gene RPL9 in Section 3.6 is not observed in SCAN-B. This discrepancy is explained in Supplement E (Xiang et al. (2023)).

4.1. Biological reproducibility of the mixture pattern. In Section 3.3 we find an interesting mixture pattern detected by the Parabolic BID in the context of the five subtypes. To investigate whether this mixture pattern is reproducible in the SCAN-B data, we chose the

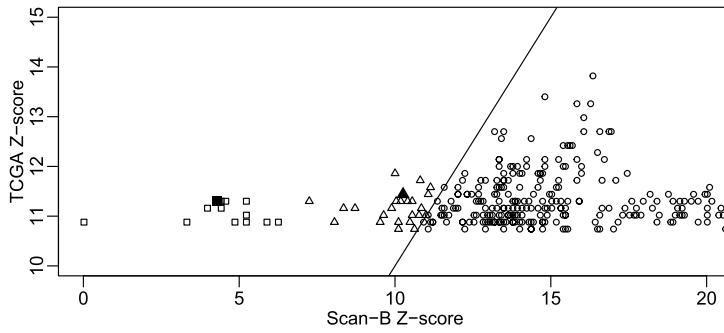


FIG. 11. Scatter plot comparing the Parabolic BID z -scores between the SCAN-B and TCGA data sets for the significant mixture pairs within the top 200 genes in the five subtypes context of TCGA. The dark square and dark triangle points are illustrated in Figures 12 and 13.

most significant 200 genes for this context in TCGA results, as shown in the network connection plot Figure 6. To understand the relationship between pairs over the two data sets, we rerun BET for these pairs in the SCAN-B data set and record the corresponding z -scores for the Parabolic BID. Thus, we compare the significance for only the mixture pattern in these two data sets, as shown in Figure 11.

Within these top 200 genes, seven genes do not exist in the SCAN-B. Considering only the remaining, 193 genes resulted in 298 significant (in TCGA) mixture pairs. The corresponding SCAN-B significance for each pair is compared in Figure 11. In particular, each point is one mixture pair whose TCGA z -score is shown on the vertical axis and SCAN-B z -score is shown on the horizontal axis.

Because the sample size is much larger for SCAN-B, stronger significance is expected for most pairs. This is highlighted using the dark line $y = x$ showing which pairs are equal. As expected, most SCAN-B z -scores are larger, reflected as circle points to the right of the line $y = x$.

Points to the left of the line $y = x$ seem to fall into two different types. For the triangle points, the TCGA z -score is not much bigger than the SCAN-B z -score, suggesting this could be just random variation. This is more carefully investigated in Figure 12. The square ones are investigated in Figure 13. These are gene pairs with substantial missing values in the SCAN-B version of the data.

Typical behavior of the triangle points is studied in Figure 12 by showing the pair represented as a dark triangle in Figure 11. The BET diagnosis (left panel) and the SCAN-B




FIG. 12. Left: BET diagnosis plot for a pair of genes in the SCAN-B data set that is shown as a dark triangle in Figure 11. Middle: The scatter plot of the same two genes in the original SCAN-B scale. Left: The scatter plot of the same two genes in the normalized TCGA count scale. This example represents gene pairs which tend to be less significant in the SCAN-B data set but still have the mixture pattern.




FIG. 13. *Left: BET diagnosis plot for a pair of genes in the SCAN-B that is shown as a dark square in Figure 11 revealing a data threshold issue. Middle: The scatter plot of the same two genes in the original SCAN-B scale. Right: The scatter plot of the same two genes in the normalized TCGA count scale. This example represents a type of gene pairs which tend to be less significant in the SCAN-B data set.*

scatter (middle panel) plots definitely show the same behavior as in Figure 5, which is a clear separation between basal and the other subtypes. However, the separation of the basal is more distinct in TCGA, as shown in the TCGA scatter plot (right panel) of the same two genes, which is consistent with the more significant TCGA z -scores. We observed similar behavior for each of the pairs represented as triangles in Figure 11.

Figure 13 shows the pair highlighted with the dark square in Figure 11, which illustrates a different phenomenon represented by the pairs symbolized by squares. The BET diagnosis (left panel) and the SCAN-B scatter (middle panel) plots reveal a data threshold issue in the SCAN-B data set. The corresponding TCGA scatter plot (right panel) of this same pair does not have this issue. This threshold effect apparently is caused by missing values in the SCAN-B data set being replaced by the minimum of their values. This same phenomenon occurred for each of the pairs represented by squares in Figure 11. As discussed in Section 3.1, we recommend handling such threshold data by jittering.

4.2. *Lack of reproducibility of the bimodal pattern.* Here we study the pair of genes that gives the strongest bimodal pattern signal in TCGA data set, which are RPL9 and RPL32, as shown in Figure 10. We rerun BET on this pair of genes in the SCAN-B data, and the strongest BID for this pair is linear, as shown in the left panel of Figure 14. Both the BET diagnosis (left panel) and the scatter (middle panel) plots show a relatively standard positively




FIG. 14. *Left: BET diagnosis plot of the most significant BID (A_1B_1) for RPL9 and RPL32 in the SCAN-B data set. Middle: The scatter plot of the same two genes in the original SCAN-B scale. Right: BET diagnosis plot of the W BID for the same two genes. This shows the bimodal pattern observed for RPL9 in TCGA data is not biologically reproducible.*

correlated linear dependence between RPL9 and RPL32. The right panel of Figure 14 gives the BET diagnosis plot for the W BID for this pair, which is much less significant than the linear pattern ($z = 30.85$ vs. $z = 10.96$). This shows that the surprising bimodal dependence observed in Figure 10 is not biologically reproducible. Instead, it seems to be a processing artifact. As noted above, deeper investigation of the artifact is given in Supplement E (Xiang et al. (2023)).

5. Conclusion. TCGA gene expression data set is an important genomics data resource that shows many dependence patterns among genes, especially some interesting nonlinear dependence patterns. We use the computationally fast and powerful dependence testing method called BET to discover significant nonlinear dependence relationships in various contexts using the breast cancer subtypes information. We find that some interesting nonlinear dependence patterns are explained biologically by the mixture of the given breast cancer subtype distributions, such as the mixture pattern for the context of five subtypes. Some relationships motivate further biological work, such as the mixture pattern for the LumA only context. We also investigate the reproducibility of these results using an independent genomics data set. This shows that the mixture pattern is reproducible while the bimodal pattern related to the gene RPL9 is not and is apparently caused by some preprocessing steps.

Acknowledgments. The results published here are in whole or part based upon data from the Cancer Genome Atlas managed by the NCI and NHGRI (dbGaP accession phs000178).

Funding. Xiang's research was supported by SAMSI, NIH/NIAMS Grants P30AR072580 and R21AR074685, and DMS-2152289 from NSF.

Perou's research was supported by NCI Breast SPORE program P50-CA58223 and U01CA238475-01.

Zhang's research was partially supported by NSF Grants DMS-1613112, IIS-1633212, DMS-1916237 and DMS-2152289.

Marron's research was supported by NSF Grants IIS-1633074 and DMS-2113404.

SUPPLEMENTARY MATERIAL

Supplement A: The relationships between top genes with respect to five nonlinear BIDs (DOI: [10.1214/23-AOAS1745SUPPA](https://doi.org/10.1214/23-AOAS1745SUPPA); .zip). A network connection plot of the relationships between the 200 most significant genes with respect to five nonlinear BIDs.

Supplement B: Gene set enrichment analysis results for the mixture dependence of the ANKS6 community (DOI: [10.1214/23-AOAS1745SUPPB](https://doi.org/10.1214/23-AOAS1745SUPPB); .zip). A full list of overlapping computation (using Gene Set Enrichment Analysis, GSEA) results of the ANKS6 Community with the gene sets in the Molecular Signatures Database (MSigDB) for the Mixture Dependence in the Five Subtypes context.

Supplement C: More examples for the gene ZDHHC2 in the LumA context for the parabolic BID (DOI: [10.1214/23-AOAS1745SUPPC](https://doi.org/10.1214/23-AOAS1745SUPPC); .zip). Two more significant parabolic BID gene pairs with ZDHHC2 in the context of only Luminal A. ZDHHC2 is on different axes in these two examples. This shows the reflection property of BET and the bifurcating role of ZDHHC2.

Supplement D: Gene set enrichment analysis results for the mixture dependence of the ZDHHC2 community (DOI: [10.1214/23-AOAS1745SUPPD](https://doi.org/10.1214/23-AOAS1745SUPPD); .zip). A full list of overlapping computation (using GSEA) results of the ZDHHC2 Community with the gene sets in the MSigDB for the Mixture Dependence in the LumA only context.

Supplement E: Deeper investigation of the bimodal structure of RPL9 (DOI: 10.1214/23-AOAS1745SUPPE; .zip). Supplement E shows that the bimodal distribution of the gene RPL9, discussed in Sections 3.6 and 4.2 is an artifact caused by the early gene mapping algorithm used to preprocess that data set.

REFERENCES

- BRUEFFER, C., VALLON-CHRISTERSSON, J., GRABAU, D., EHINGER, A., HÄKKINEN, J., HEGARDT, C., MALINA, J., CHEN, Y., BENDAHL, P.-O. et al. (2018). Clinical value of RNA sequencing-based classifiers for prediction of the five conventional breast cancer biomarkers: A report from the population-based multicenter Sweden cancerome analysis network—breast initiative. *JCO Precis. Oncol.* **2** 1–18. <https://doi.org/10.1200/PO.17.00135>
- CHARAFE-JAUFFRET, E., GINESTIER, C., MONVILLE, F., FINETTI, P., ADELAÏDE, J., CERVERA, N., FEKAIRI, S., XERRI, L., JACQUEMIER, J. et al. (2006). Gene expression profiling of breast cell lines identifies potential new basal markers. *Oncogene* **25** 2273–2284. <https://doi.org/10.1038/sj.onc.1209254>
- CIRIELLO, G., GATZA, M. L., BECK, A. H., WILKERSON, M. D., RHIE, S. K., PASTORE, A., ZHANG, H., MCLELLAN, M., YAU, C. et al. (2015). Comprehensive molecular portraits of invasive lobular breast cancer. *Cell* **163** 506–519. <https://doi.org/10.1016/j.cell.2015.09.033>
- HOEFFDING, W. (1948). A non-parametric test of independence. *Ann. Math. Stat.* **19** 546–557. [MR0029139 https://doi.org/10.1214/aoms/1177730150](https://doi.org/10.1214/aoms/1177730150)
- IGLESIA, M. D., PARKER, J. S., HOADLEY, K. A., SERODY, J. S., PEROU, C. M. and VINCENT, B. G. (2016). Genomic analysis of immune cell infiltrates across 11 tumor types. *J. Natl. Cancer Inst.* **108**. <https://doi.org/10.1093/jnci/djw144>
- KAC, M. (1959). *Statistical Independence in Probability, Analysis and Number Theory. The Carus Mathematical Monographs* **12**. Math. Assoc. of America; distributed by Wiley, New York. [MR0110114](https://doi.org/10.1093/jnci/djw144)
- KINNEY, J. B. and ATWAL, G. S. (2014). Equitability, mutual information, and the maximal information coefficient. *Proc. Natl. Acad. Sci. USA* **111** 3354–3359. [MR3200177 https://doi.org/10.1073/pnas.1309933111](https://doi.org/10.1073/pnas.1309933111)
- KRAKOV, A., STÖGBAUER, H. and GRASSBERGER, P. (2004). Estimating mutual information. *Phys. Rev. E* (3) **69** 066138, 16. [MR2096503 https://doi.org/10.1103/PhysRevE.69.066138](https://doi.org/10.1103/PhysRevE.69.066138)
- LIBERZON, A., SUBRAMANIAN, A., PINCHBACK, R., THORVALDSDÓTTIR, H., TAMAYO, P. and MESIROV, J. P. (2011). Molecular signatures database (MSigDB) 3.0. *Bioinformatics* **27** 1739–1740. <https://doi.org/10.1093/bioinformatics/btr260>
- LIBERZON, A., BIRGER, C., THORVALDSDÓTTIR, H., GHANDI, M., MESIROV, J. P. and TAMAYO, P. (2015). The Molecular Signatures Database (MSigDB) hallmark gene set collection. *Cell Syst.* **1** 417–425. <https://doi.org/10.1016/j.cels.2015.12.004>
- PARKER, J. S., MULLINS, M., CHEANG, M. C., LEUNG, S., VODUC, D., VICKERY, T., DAVIES, S., FAURON, C., HE, X. et al. (2009). Supervised risk predictor of breast cancer based on intrinsic subtypes. *J. Clin. Oncol.* **27** 1160.
- PEROU, C. M., SØRLIE, T., EISEN, M. B., VAN DE RIJN, M., JEFFREY, S. S., REES, C. A., POLLACK, J. R., ROSS, D. T., JOHNSEN, H. et al. (2000). Molecular portraits of human breast tumours. *Nature* **406** 747–752.
- SAAL, L., VALLON-CHRISTERSSON, J., HÄKKINEN, J., HEGARDT, C., GRABAU, D., WINTER, C., BRUEFFER, C., TANG, M.-H. E., REUTERSWÄRD, C. et al. (2015). The Sweden Cancerome Analysis Network—Breast (SCAN-B) initiative: A large-scale multicenter infrastructure towards implementation of breast cancer genomic analyses in the clinical routine. *Gen. Med.* **7** 20. <https://doi.org/10.1186/s13073-015-0131-9>
- SUBRAMANIAN, A., TAMAYO, P., MOOTHA, V. K., MUKHERJEE, S., EBERT, B. L., GILLETTE, M. A., PAULOVICH, A., POMEROY, S. L., GOLUB, T. R. et al. (2005). Gene set enrichment analysis: A knowledge-based approach for interpreting genome-wide expression profiles. *Proc. Natl. Acad. Sci. USA* **102** 15545–15550. <https://doi.org/10.1073/pnas.0506580102>
- SUN, N. and ZHAO, H. (2014). Putting things in order. *Proc. Natl. Acad. Sci. USA* **111** 16236–16237. <https://doi.org/10.1073/pnas.1418862111>
- SZÉKELY, G. J. and RIZZO, M. L. (2013). Energy statistics: A class of statistics based on distances. *J. Statist. Plann. Inference* **143** 1249–1272. [MR3055745 https://doi.org/10.1016/j.jspi.2013.03.018](https://doi.org/10.1016/j.jspi.2013.03.018)
- SZÉKELY, G. J., RIZZO, M. L. and BAKIROV, N. K. (2007). Measuring and testing dependence by correlation of distances. *Ann. Statist.* **35** 2769–2794. [MR2382665 https://doi.org/10.1214/009053607000000505](https://doi.org/10.1214/009053607000000505)
- THE CANCER GENOME ATLAS NETWORK (2012). Comprehensive molecular portraits of human breast tumours. *Nature* **490** 61–70.
- WILKINSON, L., ANAND, A. and GROSSMAN, R. (2005). Graph-theoretic scagnostics. In *IEEE Symposium on Information Visualization* 157–164. IEEE Comput. Soc., Los Alamitos.

- XIANG, S., ZHANG, W., LIU, S., HOADLEY, K. A., PEROU, C. M., ZHANG, K. and MAR-
RON, J. S. (2023). Supplement to “Pairwise nonlinear dependence analysis of genomic
data.” <https://doi.org/10.1214/23-AOAS1745SUPPA>, <https://doi.org/10.1214/23-AOAS1745SUPPB>,
<https://doi.org/10.1214/23-AOAS1745SUPPC>, <https://doi.org/10.1214/23-AOAS1745SUPPD>,
<https://doi.org/10.1214/23-AOAS1745SUPPE>
- ZHANG, K. (2019). BET on independence. *J. Amer. Statist. Assoc.* **114** 1620–1637. MR4047288
<https://doi.org/10.1080/01621459.2018.1537921>