

TESTING THE DISJUNCTION HYPOTHESIS USING VORONOI DIAGRAMS WITH APPLICATIONS TO GENETICS¹

BY DAISY PHILLIPS AND DEBASHIS GHOSH

Pennsylvania State University

Testing of the disjunction hypothesis is appropriate when each gene or location studied is associated with multiple p -values, each of which is of individual interest. This can occur when more than one aspect of an underlying process is measured. For example, cancer researchers may hope to detect genes that are both differentially expressed on a transcriptomic level and show evidence of copy number aberration. Currently used methods of p -value combination for this setting are overly conservative, resulting in very low power for detection. In this work, we introduce a method to test the disjunction hypothesis by using cumulative areas from the Voronoi diagram of two-dimensional vectors of p -values. Our method offers much improved power over existing methods, even in challenging situations, while maintaining appropriate error control. We apply the approach to data from two published studies: the first aims to detect periodic genes of the organism *Schizosaccharomyces pombe*, and the second aims to identify genes associated with prostate cancer.

1. Introduction. In current genetics and biological research, frequently thousands of hypothesis tests are performed simultaneously. With such a large number of tests, control of the family-wise error rate (FWER) is overly conservative, resulting in low power for detection of true alternative signals. For this reason, False Discovery Rate (FDR) control is an intensely studied topic of interest. Methods such as the Benjamini–Hochberg (B–H) procedure [Benjamini and Hochberg (1995)] control FDR when each hypothesis test is associated with a single p -value. Many refinements have been proposed to increase power [Benjamini and Hochberg (2000), Benjamini, Krieger and Yekutieli (2006), Ghosh (2011), Storey (2002)]. Pounds (2006) gives a good summary of these approaches. However, when multiple p -values are associated with each hypothesis, these methods alone are insufficient for declaring significance while controlling FDR.

Multiple p -values can be considered for each hypothesis test in settings such as repeated experiments in the same study, meta-analysis across multiple studies and measurement of multiple aspects of a single underlying process of interest. To perform a single hypothesis test using multiple p -values, techniques for p -value

Received March 2012; revised October 2013.

¹Supported from the National Science Foundation Grant (Award Number ABI-1262538).

Key words and phrases. Multiple testing, false discovery rates, Voronoi tessellations, empirical null distributions.

combination are frequently used to create a single summarized value. Multiple comparison adjustments can then be applied to these summarized values.

With this move from single p -values to vectors of p -values, which we refer to here as p -vectors, clear specification of the null and alternative hypotheses is critical. If the goal is to pool information, for example, when testing p -vectors from repeated experiments or from multiple studies, then conjunction or partial conjunction hypotheses are appropriate. The conjunction null hypothesis is that all p -values contained in a p -vector are from a null distribution, and rejection is possible when at least one p -value shows evidence of being from a nonnull distribution. Rejection of the partial conjunction hypothesis requires at least u of n p -values to show evidence of being from nonnull distributions. There are scenarios, however, when the hypothesis associated with each p -value of the p -vector is of interest individually, and rejection should be possible only when there is evidence that all such hypotheses are nonnull. In this case, the disjunction hypothesis is of primary interest. Distinctions between the conjunction, partial conjunction and disjunction hypotheses are further described in Section 2 of this paper.

Testing of the disjunction hypothesis is appropriate when multiple aspects of a single underlying biological process are measured. For example, there is interest in detection of genes related to cancer progression that are both differentially expressed on a transcriptomic level and show evidence of copy number aberrations in cancerous tissue [Fritz et al. (2002), Kim et al. (2007), Pollack et al. (2002), Tonon et al. (2005), Tsafrir et al. (2006)]. Another motivating example is detection of periodic genes as explored by de Lichtenberg et al. (2005). In this case, the disjunction hypothesis is considered through one p -value for periodicity and a second for regulation of expression. The most commonly used summary method for the disjunction hypothesis uses the maximum of all p -values for each test [Wilkinson (1951)] and typically has very low power.

This paper presents an approach for p -value combination appropriate for testing the disjunction hypothesis when there are two p -values associated with each gene or location. The approach considers p -vectors as locations on the unit square, where certain challenges absent in the case of single p -values arise. First, the strict ordering of p -values on the real line is lost. Second, relationships between p -vectors are complicated and, third, their components may be correlated. In light of these challenges, a method for large-scale simultaneous testing of the disjunction hypothesis must accomplish three objectives. It must account for the relative positioning of the p -vectors in the plane, allow for multiple ordering schemes and, finally, allow for FDR control under any correlation structure of the test statistics used to calculate the p -vectors' components.

The approach proposed here addresses these challenges through the use of Voronoi diagrams, flexible incorporation of ordering schemes and empirical null distributions [Efron (2004)]. This paper is organized as follows. Section 2 details the disjunction hypothesis framework while Section 3 describes background on control of FDR and existing univariate procedures. Sections 4 and 5 introduce

Voronoi diagrams and multiple ordering schemes for p -vectors. Section 6 describes a technique for summarizing p -vectors to a single value and details how these values can be used to control FDR. We explore properties of the procedure through simulations in Section 7. A possible extension for higher dimensional p -vectors is discussed in Section 8. We apply the procedure to two genomic studies in Sections 9 and 10.

2. The disjunction of null hypotheses. When multiple p -values are associated with a hypothesis test, the interpretation of significance depends on the specification of null and alternative hypotheses. Consider m p -vectors, each of length n , denoted

$$(1) \quad P_i = (p_{i1}, \dots, p_{in}), \quad i = 1, \dots, m.$$

In the context of a genomic study, i is the index of the individual gene, while n is the number of p -values associated with each gene. We employ notation used by [Benjamini and Heller \(2008\)](#) to describe null and alternative hypotheses. Testing the global null hypothesis, also known as the conjunction of the null hypotheses, is equivalent to testing that *at least one* of the p -values p_{i1}, \dots, p_{in} is significant:

$$H_0^{1/n}: \quad \text{all hypotheses associated with } P_i \text{ are null,}$$

$$H_A^{1/n}: \quad \text{at least one hypothesis associated with } P_i \text{ is nonnull.}$$

P -value combination methods for testing the conjunction null include the well-known Fisher's and Stouffer's methods for combining p -values [[Fisher \(1932\)](#), [Stouffer et al. \(1949\)](#)]. A comparison of these and other methods is presented by [Loughin \(2004\)](#). Rejection of the conjunction null can result from the influence of a single highly significant p -value even when all other p -values show no evidence for the alternative hypothesis. In this setting, the scientific conclusion from rejection is not as strong it would be if a level of increased consistency across p -values was enforced.

[Benjamini and Heller](#) proposed techniques for addressing this weakness through testing of the partial conjunction hypothesis:

$$H_0^{u/n}: \quad \text{at least } u \text{ of } n \text{ hypotheses associated with } P_i \text{ are null,}$$

$$H_A^{u/n}: \quad \text{at least } u \text{ of } n \text{ hypotheses associated with } P_i \text{ are nonnull.}$$

This hypothesis requires a level of consistency of evidence across u studies that is unnecessary in the conjunction framework, while still allowing lack of significance for some associated p -values. It can be considered a compromise between the conjunction and disjunction hypotheses. The disjunction hypothesis is also referred to as the disjunction of the null hypotheses and can be expressed as follows:

$$H_0^{n/n}: \quad \text{at least one hypothesis associated with } P_i \text{ is null,}$$

$$H_A^{n/n}: \quad \text{all hypotheses associated with } P_i \text{ are nonnull.}$$

TABLE 1
Outcomes for testing m hypotheses

	Fail to reject	Reject	Total
True null	U	V	m_0
True alternative	T	S	$m - m_0$
Total	$m - R$	R	m

This hypothesis is desirable when considering multiple p -values per test that are each of individual interest. The established p -value combination approach for testing the disjunction hypothesis is to simply select the maximum p -value of each p -vector [Wilkinson (1951)]. Error control procedures can then be applied to these maximum values. This approach is generally conservative and exhibits low power. The procedure described in this paper is suitable for testing the disjunction hypothesis and results in a gain of power over the maximum method of the p -value combination.

3. The False Discovery Rate and review of existing procedures. To define FDR we consider Table 1. The False Discovery Rate is defined to be the *expected proportion of false rejections to total rejections*, or $E[V/R]$ for $R > 0$, and 0 if $R = 0$. The FWER is defined to be the *probability of at least one false rejection*, $P(V > 0)$. Particularly in studies testing thousands or tens of thousands of hypotheses simultaneously, control of FDR grants additional power for rejection relative to control of FWER. Allowance of a controlled proportion of false positives enables increased detection of more true signals relative to limiting rejections based on the probability of at least one false rejection as required by FWER control.

We next provide a brief review of two existing procedures that control FDR when each hypothesis test has exactly one p -value: the Benjamini–Hochberg (B–H) procedure and the Generalized Benjamini–Hochberg procedure. The latter motivates our approach to p -value combination for two-dimensional p -vectors.

3.1. *The Benjamini–Hochberg procedure and its generalization.* Proposed by Benjamini and Hochberg (1995), the B–H procedure works as follows. Assume m continuous p -values p_1, \dots, p_m and that low values indicate evidence against the null. Order them as $p_{(1)}, \dots, p_{(m)}$ and compare to the thresholds $\alpha/m, 2\alpha/m, \dots, \alpha$. Define

$$(2) \quad \hat{k} = \max \left\{ i : p_{(i)} \leq \frac{i}{m} \alpha \right\}.$$

If the set in (2) is nonempty, reject the hypotheses associated with $p_{(1)}$ through $p_{(\hat{k})}$, otherwise reject nothing. Benjamini and Hochberg show that this procedure controls the FDR at a nominal level α .

Ghosh (2011) proposed a family of testing procedures based on the spacings of p -values. We present the basic procedure here. Again assume m independent p -values and order them as $p_{(1)}, \dots, p_{(m)}$. Define the spacings [Pyke (1965)] as

$$(3) \quad \tilde{p}_i = p_{(i)} - p_{(i-1)}, \quad i = 1, \dots, m + 1,$$

where $p_{(0)} = 0$ and $p_{(m+1)} = 1$. Under the null hypothesis the original p -values are distributed as Uniform(0, 1), whence the spacings \tilde{p}_i are marginally distributed as Beta(1, m). It is also simple to calculate certain expectations. Specifically, $E[\tilde{p}_i] = (m + 1)^{-1} = E[\tilde{p}_1]$.

Recall that the B–H procedure defines \hat{k} by comparing $p_{(1)}, p_{(2)}, \dots, p_{(m)}$ to $\alpha/m, 2\alpha/m, \dots, \alpha$. These quantities can be redefined as

$$(4) \quad p_{(i)} = \sum_{j=1}^i \tilde{p}_j \quad \text{and} \quad \frac{i}{m}\alpha = i \cdot \frac{m+1}{m} \cdot \frac{1}{m+1} \cdot \alpha = i \cdot \frac{m+1}{m} \cdot E[\tilde{p}_1] \cdot \alpha.$$

By substituting these quantities in the B–H procedure, the original \hat{k} can be rewritten in terms of the spacings as follows [Ghosh (2011, 2012)]:

$$(5) \quad \hat{k} = \max \left\{ i : \frac{1}{i} \sum_{j=1}^i \tilde{p}_j \leq \frac{m+1}{m} E[\tilde{p}_1] \cdot \alpha \right\}.$$

According to Benjamini and Hochberg, this procedure controls FDR at $\alpha m_0/m$, where m_0 is the number of p -values associated with the null hypothesis. In equation (5) there is an extra factor $(m + 1)/m$. Elimination of this factor results in a slightly more conservative procedure that preserves FDR control. Thus, \tilde{k} is defined as

$$(6) \quad \tilde{k} = \max \left\{ i : \frac{1}{i} \sum_{j=1}^i \tilde{p}_j \leq \alpha E[\tilde{p}_1] \right\}.$$

The definition of \tilde{k} in (6) hinges on a comparison of the average spacings between ordered p -values to the value $\alpha E[\tilde{p}_1]$. When there are numerous significant p -values their spacings are small in comparison to the expectation of a spacing under the null hypothesis. Detecting this change from the small spacings of significant p -values to the larger spacings of null p -values is also the motivation for the procedure described in this paper.

4. The Voronoi diagram. In the plane the spacings of two-dimensional p -vectors are more difficult to characterize. One generalization of “spacing” is the Voronoi diagram. It is a partition of the plane generated by an input set of points that creates a cell around each input consisting of the set of all points closer to that input than to any other. The basic properties are described by Okabe et al. (2010).

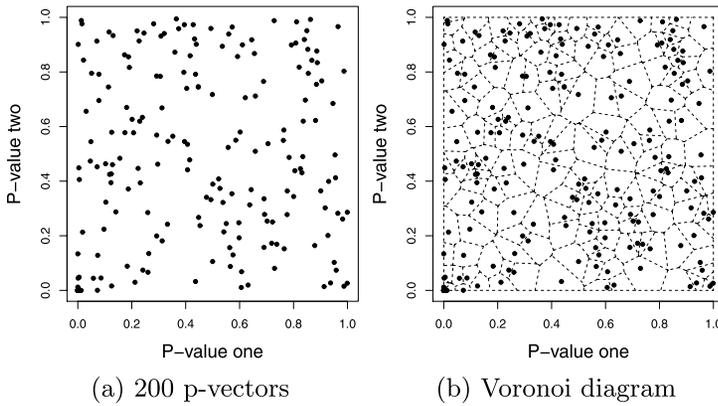


FIG. 1. (a) 200 simulated p -vectors and (b) the corresponding Voronoi diagram.

In the setting of two-dimensional p -vectors, the Voronoi diagram partitions the unit square. For each p -vector, P_i , the diagram creates a cell, C_i , consisting of all points closer to P_i than to any other p -vector. As Jiménez and Yukich (2002) discuss, Voronoi diagrams are suitable for extension of the concept of one-dimensional spacings into higher dimensions. An illustration of the diagram for a sample set of points is presented in Figure 1(b). We follow this example in Section 5 before switching to a larger sample for subsequent sections.

Voronoi cells have many desirable properties that extend spacings to the plane. Their area and shape reflect the relative positioning of the input points. For example, clusters of inputs will have smaller cell areas than uniformly distributed inputs. Similarly, if the inputs have correlated components, there will be an increased concentration of p -vectors along the diagonal of the unit square. The Voronoi cells for the inputs along this diagonal will be smaller than those near the edge of the clustering. Our procedure uses the areas of the Voronoi cells generated by the set of p -vectors to account for their relative positions. To compute each diagram and calculate the cell areas, we use the R package `deldir` developed by Turner (2013) and available through CRAN (<http://cran.r-project.org/web/packages/deldir/index.html>).

5. Multiple ordering schemes. Recall from Section 3.1 that spacings were defined as the difference between consecutive p -values. This definition is dependent upon an ordering of the p -values that is unique on the real line. However, this uniqueness is lost when p -vectors are considered as bivariate locations in the unit square. We present and test multiple ordering schemes for the plane, while continuing to assume that small values of p_{i1} and p_{i2} indicate evidence against the null. For this reason the orderings begin at the origin, and each p -vector is ranked according to increasing values of D , its distance from the origin. Each scheme defines D differently. Thus, for each definition $P_{(1)}$ is the p -vector with the smallest

value of D , and $P_{(m)}$ with the largest. Here we describe the definition of D for each scheme:

1. *Euclidean ordering* results in a movement from the origin in contours with the shape of circles. Define $D^{(E)}$ as the Euclidean distance from the origin

$$(7) \quad D_i^{(E)} = \sqrt{p_{i1}^2 + p_{i2}^2} \quad (i = 1, \dots, m).$$

2. *Maximum ordering* results in contours with the shape of squares. Define $D^{(M)}$ as

$$(8) \quad D_i^{(M)} = \max\{p_{i1}, p_{i2}\} \quad (i = 1, \dots, m).$$

3. *Summation ordering* is equivalent to beginning at the origin and moving out in contours of right isosceles triangles. In this case, $D^{(S)}$ is defined as

$$(9) \quad D_i^{(S)} = p_{i1} + p_{i2} \quad (i = 1, \dots, m).$$

4. *de Lichtenberg ordering* is a ranking scheme proposed by de Lichtenberg et al. (2005). The scheme defines

$$(10) \quad D_i^{(L)} = p_{i1}p_{i2} \left(1 + \left(\frac{p_{i1}}{0.001}\right)^2\right) \left(1 + \left(\frac{p_{i2}}{0.001}\right)^2\right) \quad (i = 1, \dots, m).$$

Note that $D^{(L)}$ consists of four multiplicative factors. The first two weight $D^{(L)}$ according to the value of each individual component, and the last two penalize p -vectors that have only one very small component. For typical p -vectors the values for $D^{(L)}$ are very large as a result of division by 0.001 of both p_{i1} and p_{i2} . This magnitude is not a concern, as the interest is only in their relative values for the purpose of ranking and the values themselves are not of particular interest. The contour lines for this ordering scheme move from the origin in lines approximating an inverse function such as $y = 1/(x^3)$.

Figure 2 illustrates these four ordering schemes using the sample set of 200 p -vectors from Section 4. Table 2 presents a numerical example using five p -vectors.

It is noteworthy that three of the four ranking schemes described have concave contour lines: Euclidean, Maximum and Summation. The remaining scheme, de Lichtenberg, has convex contour lines. As we will see, these characteristics have important implications for error control.

6. Summarizing p -vectors and declaring significance. The described ranking schemes can be combined with Voronoi cell areas to summarize each ranked p -vector as a single value in the interval $(0, 1)$. Define $A_{(i)}$ as the area of the Voronoi cell associated with the ordered p -vector $P_{(i)}$ and the cumulative sum of these ordered areas as

$$(11) \quad T_{(i)} = \sum_{j=1}^i A_{(j)} \quad (i = 1, \dots, m).$$

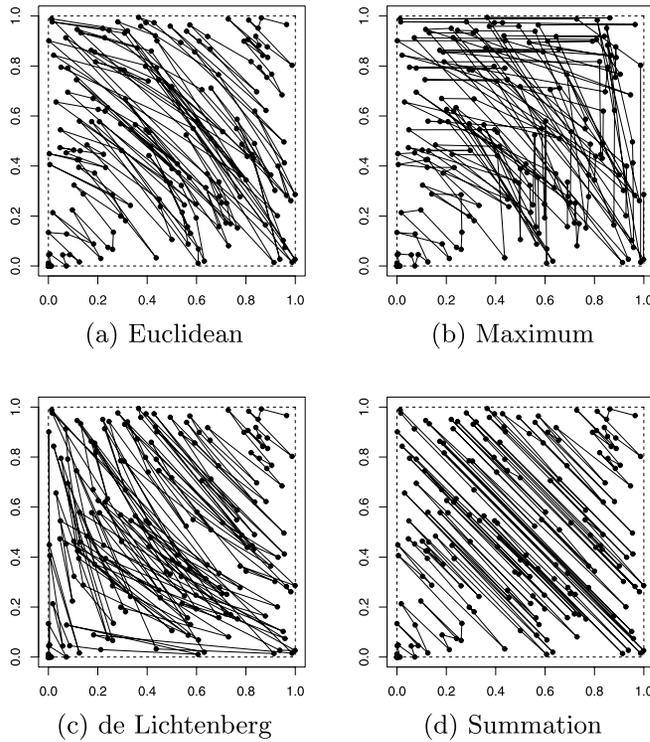


FIG. 2. Illustration of all four ordering schemes with the sample set of 200 p -vectors. The solid lines join p -vectors that are consecutively ranked, that is, $P_{(1)}$ to $P_{(2)}$, $P_{(2)}$ to $P_{(3)}$, etc.

These cumulative sums serve as combined p -values in an analogous manner that cumulative spacings comprise p -values in one dimension. These $T_{(i)}$ reflect both the relative positioning of the p -vectors in space and their distance from the origin. They can be used to make decisions in the hypothesis testing framework. Figure 3 illustrates a sample set of 1000 p -vectors and a histogram of their cumulative

TABLE 2
 Example of ordering results for five p -vectors. For each ranking scheme the distance, D , is presented for each p -vector along with the resulting rank in parentheses

P_i	$D_i^{(E)}$		$D_i^{(M)}$		$D_i^{(S)}$		$D_i^{(L)}$	
(0.85, 0.51)	0.99	(3)	0.85	(3)	1.36	(4)	$8.1 \cdot 10^{10}$	(4)
(0.91, 0.80)	1.21	(5)	0.91	(4)	1.71	(5)	$3.9 \cdot 10^{11}$	(5)
(0.23, 0.97)	1.00	(4)	0.97	(5)	1.20	(3)	$1.1 \cdot 10^{10}$	(3)
(0.62, 0.34)	0.71	(2)	0.62	(1)	0.96	(2)	$9.5 \cdot 10^9$	(2)
(0.07, 0.63)	0.63	(1)	0.63	(2)	0.79	(1)	$8.9 \cdot 10^7$	(1)

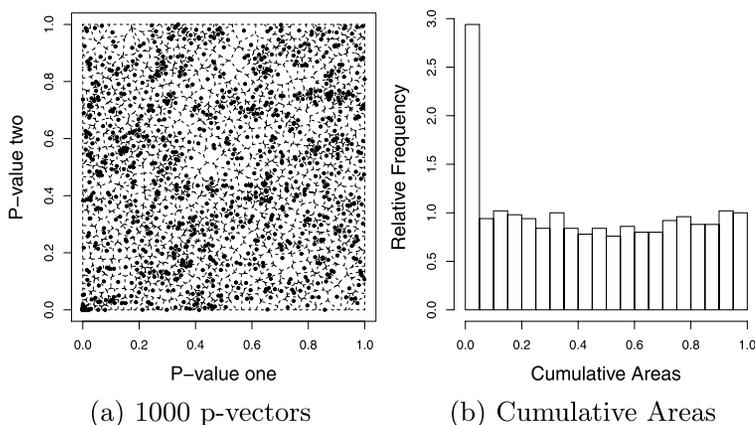


FIG. 3. An example of (a) 1000 simulated p -vectors with independent components and (b) a histogram of their cumulative areas calculated using the Euclidean ordering scheme. The sharp spike in the histogram corresponds to the p -vectors associated with alternative hypotheses.

areas. In this example the components of the p -vectors are independent, and 10% of p -vectors are associated with an alternative hypothesis.

6.1. *Multiple hypothesis testing under independence.* When the components of the p -vectors are assumed to be independent, standard multiple comparisons procedures such as B–H can be applied to the cumulative spacings with very good results. Simulation studies were performed to test the properties of FDR control and the power of this method. For each study 100 sets of test statistics were generated according to

$$(t_{i1}, t_{i2}) \sim \text{MVN} \left(\begin{pmatrix} \mu_i \\ \mu_i \end{pmatrix}, \begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix} \right), \quad i = 1, \dots, 2000;$$

10% of statistics were associated with an alternative hypothesis ($\mu_i = \mu_A$), while the remaining 90% were null ($\mu_i = 0$). P -vectors were formed from 2-sided p -values according to

$$P_i = (p_{i1}, p_{i2}) = (P(Z > |t_{i1}|), P(Z > |t_{i2}|)) \quad \text{where } Z \sim N(0, 1)$$

for $i = 1, \dots, 2000$. The proposed method was applied to each data set, using the four described ordering schemes from Section 5. Additionally, the existing p -value combination technique based on the maximum was applied. After applying the B–H procedure to each set of summary values, the FDR and 1-non discovery rate (NDR) was recorded. Using notation from Table 1, 1-NDR is defined as $E[S/(S+T)]$. This quantity can be viewed as a measure of power.

Tables 3 and 4 summarize the results for studies where $\mu_A = 2, 3, 4$, respectively. The results show that, under all ordering schemes, the proposed combination method results in greatly increased power. All concave schemes (Euclidean,

TABLE 3
Power (1-NDR) results of simulation studies under independence

	Euclidean	Maximum	Summation	de Lichtenberg	Existing procedure
$\mu_A = 2$	0.200	0.189	0.216	0.205	0.005
$\mu_A = 3$	0.772	0.760	0.788	0.796	0.098
$\mu_A = 4$	0.976	0.975	0.977	0.979	0.744

Maximum and Summation) control FDR at the desired level $\alpha = 0.05$, but the convex de Lichtenberg scheme does not. This difference becomes more pronounced when p -vectors with correlated components are considered. Additional simulations using correlated test statistics show that application of the B–H procedure to combined values is insufficient to control FDR when the correlation between test statistics surpasses 0.2. This loss of FDR control is a result of the increased concentration of p -vectors along the diagonal of the unit square under correlation, which changes the characteristics of the cumulative areas. In Section 6.2 we discuss approaches appropriate for multiple testing in these conditions.

6.2. *Multiple hypothesis testing under dependence.* In certain settings the individual components of p -vectors may be correlated. For example, correlation between components may occur when different but related aspects of an underlying biological process are measured. Any technique used for testing the disjunction hypothesis in this setting should be robust to this structure. Using an empirical null approach [Efron (2004, 2007)] in the place of the B–H procedure results in FDR control for all positive correlation structures, although the trade-off is decreased power in the case of independent components.

The use of an empirical null for determining statistical significance was proposed by Efron (2004). We consider a transformation of the summarized cumulative areas $T_{(i)}$ as defined in (11):

$$(12) \quad Z_{(i)} = \Phi^{-1}(T_{(i)}) \quad (i = 1, \dots, m),$$

where Φ is the cumulative distribution function for the standard normal random variable. In the empirical null framework, these transformed values are assumed

TABLE 4
False Discovery Rate results of simulation studies under independence

	Euclidean	Maximum	Summation	de Lichtenberg	Existing procedure
$\mu_A = 2$	0.041	0.037	0.048	0.041	0.000
$\mu_A = 3$	0.042	0.038	0.049	0.056	0.000
$\mu_A = 4$	0.042	0.040	0.045	0.053	0.000

to be from a mixture distribution

$$(13) \quad Z_{(i)} \sim f(z) = \delta f_0(z) + (1 - \delta) f_1(z) \quad (i = 1, \dots, m),$$

where f_0 is the null, or “uninteresting” distribution, and f_1 is the alternative, or “interesting” distribution. Under the theoretical null hypothesis, f_0 is the $N(0, 1)$ distribution, however, in large-scale multiple testing problems the majority of values may behave differently. When this is the case, use of an empirically determined f_0 in place of $N(0, 1)$ has important implications for the resulting inference. For example, if f_0 is estimated to be $N(0, 1.3)$, then inference based on the assumptions of $N(0, 1)$ would result in elevated type I error. Similarly, if the empirical null has smaller variance, then its use results in a gain of power without sacrificing type I error control.

Figure 4 illustrates the effect of the transformation from $T_{(i)}$ to $Z_{(i)}$ for a sample set of 1000 p -vectors. The transformation makes it much easier to detect deviations from the null hypothesis, as true alternative hypotheses are presented as a second, smaller peak to the left of the null distribution instead of in a single spike for the original values. The bivariate test statistics used to calculate these p -vectors had a correlation of 0.7. The histogram of transformed values in Figure 4(c) shows evidence of a null distribution that differs from $N(0, 1)$, as the dependence structure of the p -vectors results in thicker tails than the theoretical null predicts. For this reason, it is desirable to use an empirical null as a basis for our inference when the components of the p -vectors show evidence of correlation.

Given the choice of an empirical null, there are several options for declaring significance while controlling FDR. Efron (2004) defined the *local false discovery rate* (fdr) as the posterior probability of a value z being from the null distribution, given the value of z :

$$(14) \quad \text{fdr}(z) = P(z_i \sim f_0 | z_i = z).$$

Another important probability that can be estimated using an empirical null is the left-tail False Discovery Rate. For each value z the corresponding left-tail FDR is defined as

$$(15) \quad \text{FDR}(z) = P(z_i \sim f_0 | z_i \leq z).$$

Inference can be made based on estimated fdr or FDR values. A variety of approaches have been developed for estimation of empirical null distributions and related values [Efron (2004), Jin and Cai (2007), Muralidharan (2010), Pounds and Morris (2003), Strimmer (2008)]. We use an R package (mixFdr) developed by Muralidharan (2010) and available through CRAN (<http://cran.r-project.org/web/packages/mixfdr>), which uses an empirical Bayes mixture method to fit an empirical null, estimate effect sizes, fdr and FDR. The use of other packages or techniques is certainly possible. The function we used in simulation studies, mixFdr, includes two tuning parameters: J , the number of distributions to be estimated,

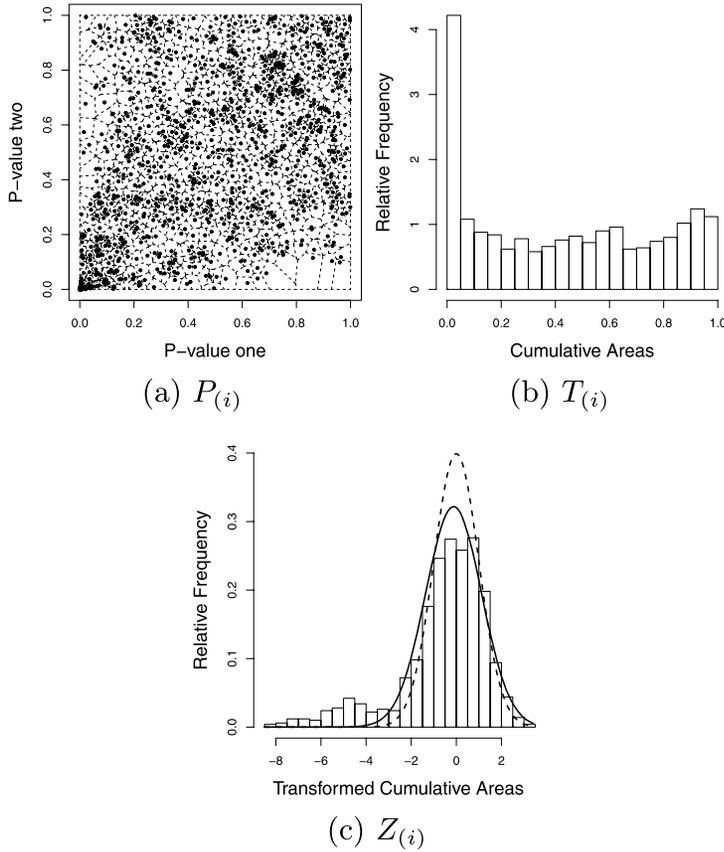


FIG. 4. An example of (a) 1000 simulated p -vectors with dependent components, (b) a histogram of their cumulative areas calculated using the Euclidean ordering scheme, and (c) a histogram of the transformed cumulative areas with empirical (solid line) and theoretical (dashed line) null distributions.

and P , a penalization parameter. A higher value of P encourages estimation of a larger null group and closer estimation of the central peak.

Careful calibration of J and P , and even experimentation with other techniques for empirical null estimation, are desirable when a single data set is under consideration. The function `mixFdr` estimates the left-tail False Discovery Rate for each $Z_{(i)}$, and we declare significant all p -vectors with these estimates of left-tail FDR less than 0.05. This approach results in appropriate error control that is robust to correlation in the components of p -vectors. Section 7 describes a simulation study performed to explore properties of power and FDR control when this approach is applied to the cumulative areas from ordered p -vectors.

7. Simulation study when components are correlated. To illustrate properties of the procedure, we ran three simulation studies: one each for strong, moder-

ate and weak alternative signals. We were interested in evaluating FDR control and power. For each simulated data set we set $m = 2000$ and generated test statistics by

$$(t_{i1}, t_{i2}) \sim \text{MVN}\left(\begin{pmatrix} \mu_i \\ \mu_i \end{pmatrix}, \begin{pmatrix} 1 & \rho \\ \rho & 1 \end{pmatrix}\right), \quad i = 1, \dots, m$$

for strong, moderate and weak alternative signals $\mu_A = 4, 3, 2$, respectively. For null test statistics, $\mu_i = 0$. 10% of test statistics for each data set were generated from the alternative distribution. P -vectors were formed from 2-sided p -values.

For each simulation study ρ varied from 0 to 0.8 in increments of 0.1, and 100 data sets were simulated for each value of ρ . We performed the procedure using all four ordering schemes on each simulated data set, using `mixFdr` to estimate empirical null distributions and left-tail FDR for each data set, rejecting all hypotheses associated with p -vectors whose estimated left-tail FDR was less than 0.05. We set $J = 2$ for all data sets, and after calibrating the fit of several example empirical nulls for weak, moderate and strong signals, we set $P = 400, 800$ and 1000 for the respective simulations. Additionally, the B–H procedure was applied to the maximum values from each p -vector to compare the proposed method to an existing approach. Figure 5 summarizes the results of all three simulation studies. We include the R code used to perform these simulations as a supplementary file [Phillips and Ghosh (2014a)].

The proposed technique for combining p -values has improved power when compared to the existing procedure. This improvement is greatest for weak and moderate alternatives. These simulations further show that ordering scheme matters, although the differences between the three convex (Euclidean, Maximum, Summation) ordering schemes are small compared to the difference between them and the de Lichtenberg ordering. The de Lichtenberg ordering displays characteristics that differ considerably from the other three. Specifically, it shows a tendency to lose FDR control as the components of the p -vectors become more correlated. We present simulation results for this ordering for the sake of completeness, but we do not recommend using it for data analysis when testing the disjunction null. The nature of its contour lines suggests this ordering scheme is in fact more appropriate for testing the conjunction or partial conjunction hypothesis, as these lines resemble the contour lines for Fisher's or Stouffer's p -value combination techniques from Figure 1 of Owen (2009).

Further simulations using the empirical null allowed evaluation of the proposed approach in the presence of p -vectors constructed from test statistics with means $(0, \mu_A)$ or $(\mu_A, 0)$. These p -vectors should not be found significant in the disjunction setting, but should be under the conjunction framework. In the presence of up to 20% of such vectors, and for correlation structures from $\rho = 0$ to $\rho = 0.8$, FDR control was maintained below the nominal 0.05 level, and power properties showed a clear advantage over existing methods. Detailed descriptions and results are included in supplementary materials [Phillips and Ghosh (2014b)].

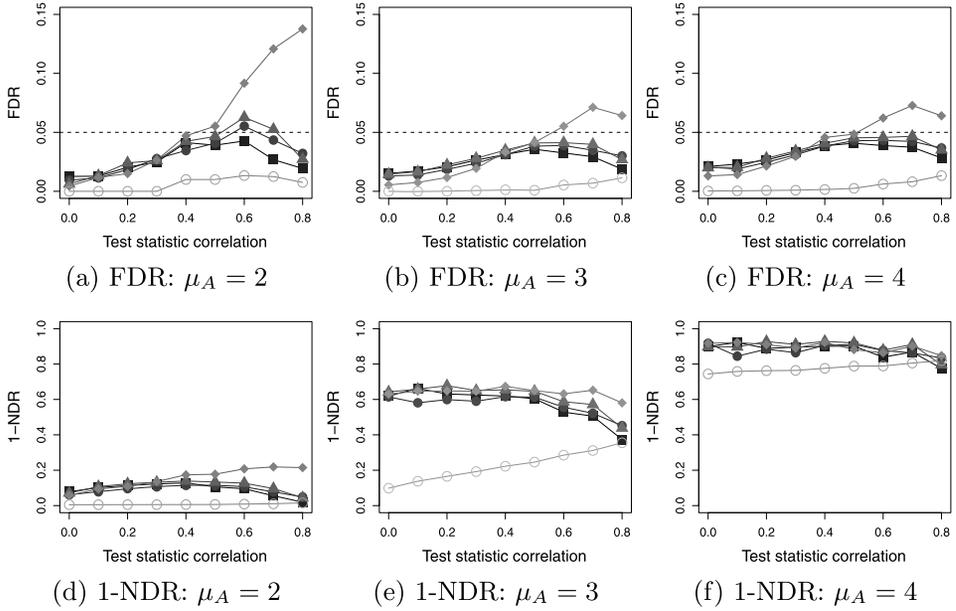


FIG. 5. Summarized results of simulation studies for test statistics with varying correlation structure. (a), (b) and (c) present FDR when alternative signals are 2, 3 and 4 standard deviations from the null. (d), (e) and (f) present 1-NDR for the same data sets. Solid squares, circles, triangles and diamonds represent Euclidean, Maximum, Summation and de Lichtenberg ordering schemes, respectively. The open circles represent the existing approach using the maximum of each component for inference.

8. Extension to higher dimensions. The approach described in this paper is suitable when there are two p -values associated with each hypothesis test, however, in many situations three or more p -values will be available. In theory, the procedure can be extended to higher dimensions by replacing cumulative areas with cumulative volumes, hypervolumes, etc. In practice, however, the computation complexity for Voronoi cells increases quickly with dimension. Average time complexity is as low as $\mathcal{O}(n)$ in the plane, but is at least $\mathcal{O}(n^2)$ in 3-space [Okabe et al. (2010)]. To avoid this disadvantage, we consider an alternative extension using the sets of all possible pairs of components. Consider a set of m 3-dimensional p -vectors:

$$(16) \quad P_i = (p_{i1}, p_{i2}, p_{i3}), \quad i = 1, \dots, m.$$

Then define three sets of 2-dimensional p -vectors constructed via a pairwise combination of components of P_i :

$$(17) \quad \{(p_{i1}, p_{i2})\}, \{(p_{i1}, p_{i3})\}, \{(p_{i2}, p_{i3})\}, \quad i = 1, \dots, m.$$

For each of these sets of two-dimensional p -vectors the Voronoi diagram is computed and cell areas saved. Thus, each p -vector P_i is associated with three

TABLE 5
Simulation results for proposed extension

	Proposed extension			Existing approach		
	$\mu_A = 2$	$\mu_A = 3$	$\mu_A = 4$	$\mu_A = 2$	$\mu_A = 3$	$\mu_A = 4$
FDR	0.023	0.004	0.023	0.000	0.000	0.000
1-NDR	0.098	0.730	0.986	0.005	0.007	0.610

individual cell areas, $A_i^{1,2}$, $A_i^{1,3}$ and $A_i^{2,3}$, as well as an average area $\bar{A}_i = (A_i^{1,2} + A_i^{1,3} + A_i^{2,3})/3$. This average area can then be used in conjunction with an ordering scheme to create the summarized areas used for inference. Define $P_{(1)}, P_{(2)}, \dots, P_{(m)}$ to be the p -vectors ranked according to a specified ordering scheme such as Euclidean distance from the origin, and $\bar{A}_{(1)}, \bar{A}_{(2)}, \dots, \bar{A}_{(m)}$ to be the corresponding average areas. Then the cumulative average areas are defined as $\bar{T}_{(i)} = \sum_{j=1}^i \bar{A}_{(j)}$.

Multiple testing can then be performed on these summarized cumulative average areas using the methods described in Sections 6.1 and 6.2. Further investigation into the properties of this approach is necessary, as well as research on other possible extensions for higher dimensions. A preliminary simulation study using three-dimensional p -vectors with independent components was conducted with weak, moderate and strong alternative test statistics. For each data set, test statistics were generated according to

$$(t_{i1}, t_{i2}, t_{i3}) \sim \text{MVN} \left(\begin{pmatrix} \mu_i \\ \mu_i \\ \mu_i \end{pmatrix}, \begin{pmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{pmatrix} \right), \quad i = 1, \dots, m.$$

Three-dimensional P -vectors were formed from 2-sided p -values. The resulting p -vectors were ordered according to Euclidean distance from the origin. Hypothesis testing was performed using the B–H procedure on the summarized cumulative average areas. The existing technique of applying the B–H procedure to the set of maximum p -values from each p -vector was also performed for comparison. Table 5 summarizes the findings of the simulations.

9. Application to *Schizosaccharomyces pombe* data. In 2004 and 2005, three papers were published investigating the periodicity of genes in the fission yeast cell *Schizosaccharomyces pombe*. Specifically, Oliva et al. (2005) produced three data sets including time points for three complete cell cycles using two different synchronization techniques. In their paper they identified 750 genes determined to be periodically expressed based on a ranking scheme and cutoff. We apply our approach to test for periodicity in a hypothesis testing framework using Fisher's exact G statistic to measure evidence of periodicity.

TABLE 6
Summary of results from Oliva et al. data sets considered separately

	Elutriation a	Elutriation b	Cdc25
Complete genes	3050	2394	3724
Evenly spaced time points	33	31	51
Genes (%) with p -values < 0.05 ,	868 (22.8%)	546 (26.4%)	2458 (66.0%)
Significant genes (%) using B–H	527 (17.3%)	155 (6.5%)	2252 (60.5%)

9.1. *The data.* Three microarray data sets from *Schizosaccharomyces pombe* from Oliva et al. were used: Elutriation a, Elutriation b and Cdc25. The first two were produced using Elutriation synchronization, and the last using a Cdc25 block-release synchronization technique. We apply our technique on the two Elutriation sets, using Fisher’s exact G statistic to calculate the p -vector for each gene. This test statistic requires evenly spaced time points, necessitating omission of any measurements that occur at uneven intervals. The Elutriation a data set includes 50 time points, however, only 33 are at regular intervals of 8 minutes. For Elutriation b, many of the time points are technical repeats. We keep the first measurement in each case, leaving 31 evenly spaced time points for each gene taken at intervals of 10 minutes. The Cdc25 data set has a total of 51 evenly spaced time points, taken at intervals of 15 minutes. Only genes with complete measurements for all selected time points are considered.

9.2. *Results using existing procedures.* The data show evidence of widespread periodicity. Considered separately, Elutriation a, Elutriation b and Cdc25 have 22.8%, 26.4% and 66% of p -values less than 0.05. Even controlling FDR using the B–H procedure on each set independently results in a very high rate of rejection. Table 6 presents a summary of the data and marginal analysis of all three data sets. Figure 6 presents histograms of the p -values when the data sets are considered independently.

Consider the p -vectors formed using p -values generated by Elutriation a data and Elutriation b data. Note that these two Elutriation data sets were both generated using the same synchronization technique, and the p -values generated by each repetition have roughly comparable marginal distributions. To test the disjunction hypothesis for Elutriation a and Elutriation b, using an existing technique the maximum p -value for each gene is preserved. The B–H procedure is then applied to these maximum values. The resulting number of rejections is 15, which is surprisingly low. Figure 7(a) helps to explain this result. The p -vectors’ components do not show evidence of correlation, thus considering only the maximum of each p -vector’s components gives a distribution that is very different from either of the marginal distributions. Our proposed approach uses information from both p -values and gives a different result.

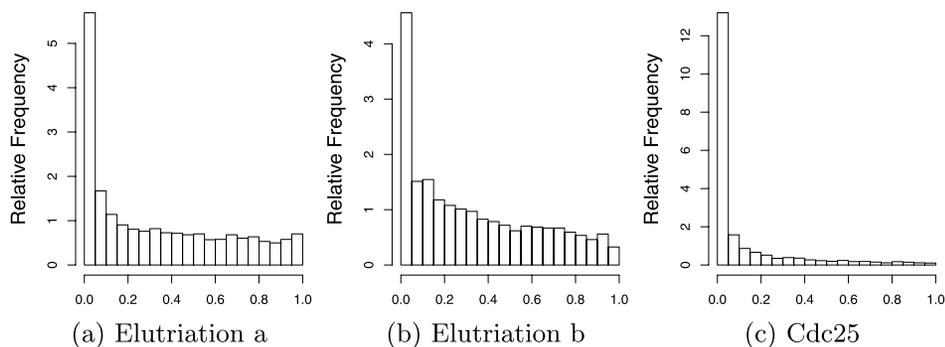


FIG. 6. Histogram of p -values for (a) Elutriation a, (b) Elutriation b and (c) Cdc25 block release. Note the strong evidence of periodicity in all three experiments, particularly Cdc25.

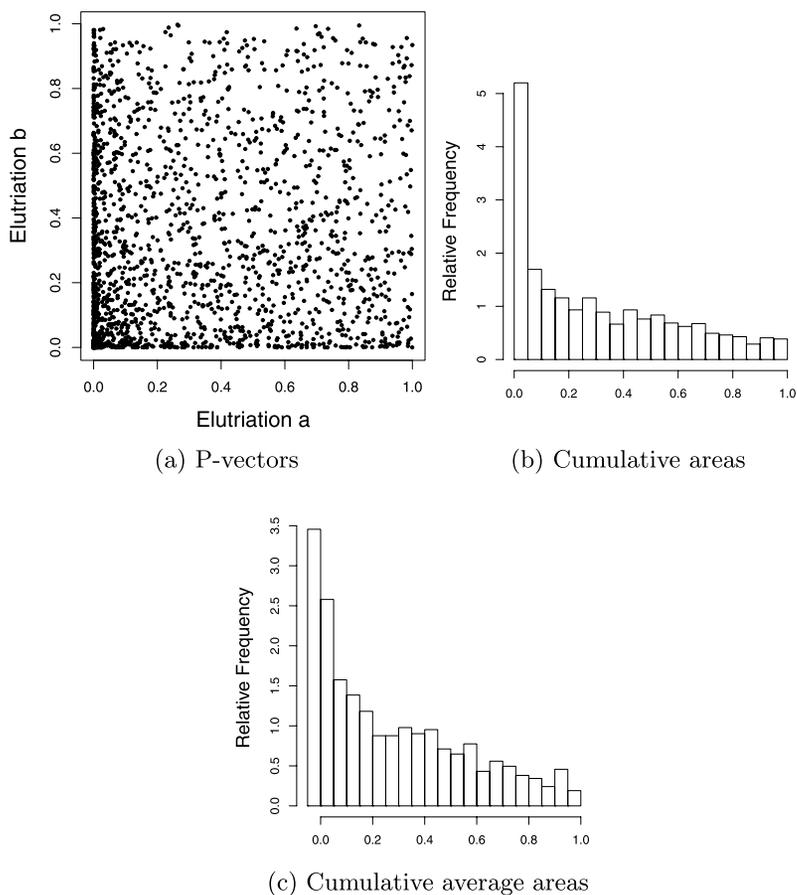


FIG. 7. (a) P -vectors formed from Fisher's G statistic of Elutriation a and Elutriation b, (b) a histogram of cumulative cell areas formed using the Euclidean ordering scheme, and (c) a histogram of cumulative average cell areas when all three experiments are considered.

9.3. *Results using Voronoi p -value combination on Elutriation data.* We apply our p -value combination method using the Euclidean, Maximum and Summation ordering schemes to the p -vectors formed from the Elutriation a and b experiments. The p -vectors are plotted in Figure 7(a). The components of the p -vectors do not show evidence of high correlation, and we apply the B–H procedure to the cumulative areas generated from each ordering scheme. This application results in 225, 213 and 249 rejections of the disjunction hypothesis using Euclidean, Maximum and Summation orderings, respectively.

Application of an empirical null approach to the combined areas yields a very different result. Because of the high amount of periodicity detected in the experiments, the empirical null is estimated to have a negative mean. This shift to the left of up to -0.87 results in rejection of far fewer genes: 15, 12 and 11 for the three concave ordering schemes. These genes could be considered significantly more periodic than the rest, although other genes also show evidence of periodicity.

The two considerations of the combined values reflect two different scientific questions. By using the B–H procedure on the combined areas, the genes found are those that show significant periodic expression in both elutriation experiments. The genes found using the empirical null procedure are those genes that are significantly periodic in both experiments relative to the majority of genes. This distinction explains the difference in numbers of genes found significant.

9.4. *Extension of procedure to include Cdc25.* The extension to three dimensions described in Section 8 can be applied to the 3-dimensional p -vectors formed from Oliva et al. data. We order the three-dimensional p -vectors according to their Euclidean distance from the origin and calculate the three Voronoi cell areas associated with each p -vector. From these we calculate each p -vector's average cell area and then cumulative average areas. Figure 7(c) presents a histogram of these values. Note that many of these cumulative areas are quite small as a result of the high number of very small p -values from the Cdc25 experiment.

Application of the B–H procedure to the cumulative average areas formed using the Euclidean ordering scheme results in rejection of 165 disjunction hypotheses. These 165 genes are those that show significant evidence of periodic expression in all three of the experiments performed by Oliva et al. The existing procedure using the maximum values yields a mere 12 rejections for these experiments. Using an empirical null approach on the transformed cumulative average areas yields results similar to those discussed in Section 9.3. Because of the evidence of widespread periodicity throughout the experiment, only 8 genes show behavior that is significantly more periodic in comparison to the majority of genes when all three experiments are considered.

10. An application related to prostate cancer. Identification of genes implicated in cancer progression is a research topic of great interest. Several studies

have shown interest in identifying genes that show both alterations in copy number and evidence of differential expression in cancerous tumors [Fritz et al. (2002), Kim et al. (2007), Pollack et al. (2002), Tonon et al. (2005), Tsafrir et al. (2006)]. We applied our method to data produced by Kim et al. in a study on prostate cancer progression. Data on copy number and gene expression was gathered for 7534 genes using prostate cell populations from low-grade and high-grade samples of cancerous tissue. For details on data acquisition and cleaning see the Kim et al. paper.

We calculated t -statistics for genetic expression and copy number aberrations comparing tissue types. For each of 7534 genes we compute a two-dimensional p -vector from the resulting 2-sided p -values based on the t -statistics. Figures 8(a) and 8(b) present histograms of the expression and copy number p -values, while Figure 8(c) presents a representation of the resulting p -vectors. Upon close inspection, it is revealed that the smallest copy number p -values are much smaller than the smallest gene expression p -values. Thus, application of the B-H proce-

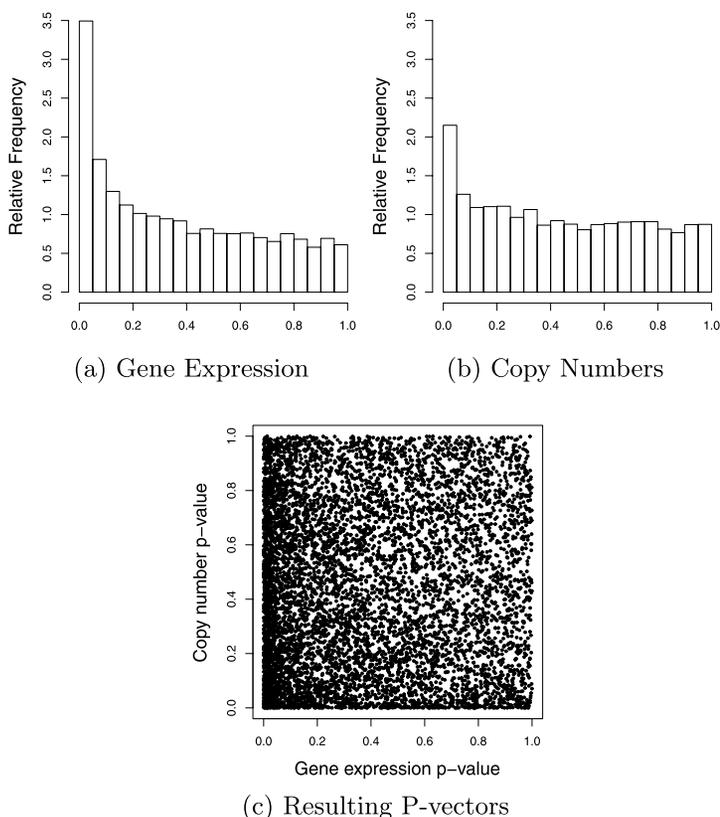


FIG. 8. Histogram of p -values for (a) expression and (b) copy number. (c) presents the resulting p -vectors in the unit square.

TABLE 7
Summary of Gene Functional Classification from DAVID

Number of genes	Enrichment score	Keywords
16	0.73	Peptidase; Serine; Endopeptidase; Kringle
11	0.58	Transmembrane; Membrane; Extracellular; Cytoplasmic
3	0.50	Transport: protein, intracellular, vesicle; Golgi apparatus
3	0.36	Catabolic process; Proteosome; Proteolysis
12	0.27	Lumen: nuclear, intracellular, organelle, membrane-enclosed; Phosphoroprotein; nucleolus; ATP binding
3	0.22	GTP-binding; Nucleotide binding: guanyl, purine; Ribonucleotide binding: guanyl, purine
20	0.12	Transmembrane; Membrane; Glycoprotein;

ture to copy number p -values yields 62 significant genes, while application to the expression p -values fails to yield any. It is unsurprising then that application of the B–H procedure to the set of all maximum p -values for each gene also produces no significant results.

Using the Voronoi p -value combination followed by the B–H procedure on the summarized values at $\alpha = 0.05$ gives 12, 14 and 25 rejections for Euclidean, Maximum and Summation orderings. Guided by the results of simulation, we consider the rejections made using the Summation ordering. Of these 25, four were mapped to official gene names, and all four were listed in the COSMIC database of cancer genes [Forbes et al. (2011)]. These four genes are CABLES2, PAK1IP1, CAMKV and TSHZ1. The COSMIC results suggest that there are mutations in these four genes that are found in a variety of cancers, thus strengthening the evidence of these genes being putative oncogenes in prostate cancer.

To use DAVID [Huang, Sherman and Lempicki (2008, 2009)] for further investigation of our results, a larger gene list was necessary. For this purpose, we performed the B–H procedure on the combined p -values at $\alpha = 0.20$. Under the summation ordering, this yielded 306 rejections, 102 of which could be mapped to recognized genes by DAVID. The functional annotation tool found significant enrichment (adjusted p -value of 0.019) in the Fibrinolysis pathway. Fibrinolysis has been associated with prostate cancer for decades [Tagnon, Whitmore and Shulman (1952)]. Tumor classifications for different malignancies have been proposed based on the behavior of this pathway [Zacharski et al. (1992)]. Results for functional classification of the 102 genes are summarized in Table 7.

11. Discussion. In this paper we have presented a novel approach to p -value combination for testing the disjunction hypothesis when two p -values are considered for each test. The approach uses an extension of one-dimensional spacings,

Voronoi cell areas, in combination with concave ordering schemes to define cumulative areas suitable for multiple testing techniques. When the majority of p -vectors have independent components, techniques such as the B–H procedure can be directly applied. If the components are correlated, empirical null techniques are more suitable. Simulation studies showed that the approach has appropriate error control properties and results in a gain of power over the existing method. This increased power is of particular interest for detection of genes related to biological processes or implicated in cancer progression.

Four candidate ordering schemes were described, and simulations were used to test their performance in several settings. The concave up ordering proposed by de Lichtenberg et al. (2005) failed to control FDR in the paradigm of the disjunction hypothesis. As discussed in Section 6.1, we suspect that concavity of an ordering's contour lines is vital to its FDR control characteristics. Specifically, as contours become increasingly concave down, the procedure is more conservative. The reverse applies when considering concave up schemes. For this reason, we recommend using the summation ordering in practice, as it represents the boundary case between concave up and down. This offers the least conservative, and thus most powerful, procedure that retains appropriate FDR control.

This approach can be extended in several meaningful directions. The conjunction or partial conjunction hypotheses could be tested by defining suitable ordering schemes such as the minimum, or product. Extension to higher dimensions is also of utmost interest, particularly considering the scale of current biological and genomic experiments. In Section 8 we described a potential extension to three or more dimensions, but further investigation of this and other techniques is necessary.

Acknowledgments. The authors would like to thank the area Editor and an anonymous referee whose comments greatly improved the quality of this paper.

SUPPLEMENTARY MATERIAL

Supplement A: Summarized results of additional simulation studies (DOI: [10.1214/13-AOAS707SUPPA](https://doi.org/10.1214/13-AOAS707SUPPA); .pdf). We present the summarized results of the proposed procedure's performance in the challenging situations described in Section 7. Results include estimated FDR and 1-NDR for each of the two settings.

Supplement B: Supplementary code and data (DOI: [10.1214/13-AOAS707SUPPB](https://doi.org/10.1214/13-AOAS707SUPPB); .zip). R code including the functions required to perform the procedure described in this paper, to replicate the described simulation studies and to perform the described data analysis. The relevant data sets are also included.

REFERENCES

BENJAMINI, Y. and HELLER, R. (2008). Screening for partial conjunction hypotheses. *Biometrics* **64** 1215–1222. [MR2522270](https://doi.org/10.1111/j.1541-0420.2008.01222.x)

- BENJAMINI, Y. and HOCHBERG, Y. (1995). Controlling the false discovery rate: A practical and powerful approach to multiple testing. *J. R. Stat. Soc. Ser. B Stat. Methodol.* **57** 289–300. [MR1325392](#)
- BENJAMINI, Y. and HOCHBERG, Y. (2000). On the adaptive control of the false discovery rate in multiple testing with independent statistics. *J. Educ. Behav. Statist.* **25** 60–83.
- BENJAMINI, Y., KRIEGER, A. M. and YEKUTIELI, D. (2006). Adaptive linear step-up procedures that control the false discovery rate. *Biometrika* **93** 491–507. [MR2261438](#)
- DE LICHTENBERG, U., JENSEN, L. J., FAUSBØLL, A., JENSEN, T. S., BORK, P. and BRUNAK, S. (2005). Comparison of computational methods for the identification of cell cycle-regulated genes. *Bioinformatics* **21** 1164–1171.
- EFRON, B. (2004). Large-scale simultaneous hypothesis testing: The choice of a null hypothesis. *J. Amer. Statist. Assoc.* **99** 96–104. [MR2054289](#)
- EFRON, B. (2007). Correlation and large-scale simultaneous significance testing. *J. Amer. Statist. Assoc.* **102** 93–103. [MR2293302](#)
- FISHER, S. R. A. (1932). *Statistical Methods for Research Workers*, 4th ed. Oliver & Boyd, Edinburgh.
- FORBES, S. A., BINDAL, N., BAMFORD, S., COLE, C., KOK, C. Y., BEARE, D., JIA, M., SHEPHERD, R., LEUNG, K., MENZIES, A. et al. (2011). COSMIC: Mining complete cancer genomes in the catalogue of somatic mutations in cancer. *Nucleic Acids Res.* **39** D945–D950.
- FRITZ, B., SCHUBERT, F., WROBEL, G., SCHWAENEN, C., WESSENDORF, S., NESSLING, M., KORZ, C., RIEKER, R. J., MONTGOMERY, K., KUCHERLAPATI, R. et al. (2002). Microarray-based copy number and expression profiling in dedifferentiated and pleomorphic liposarcoma. *Cancer Res.* **62** 2993–2998.
- GHOSH, D. (2011). Generalized Benjamini–Hochberg procedures using spacings. *J. Indian Soc. Agricultural Statist.* **65** 213–220, 262. [MR2868128](#)
- GHOSH, D. (2012). Incorporating the empirical null hypothesis into the Benjamini–Hochberg procedure. *Stat. Appl. Genet. Mol. Biol.* **11** Art. 11, front matter+19. [MR2958610](#)
- HUANG, D. W., SHERMAN, B. T. and LEMPICKI, R. A. (2008). Systematic and integrative analysis of large gene lists using DAVID bioinformatics resources. *Nature Protoc.* **4** 44–57.
- HUANG, D. W., SHERMAN, B. T. and LEMPICKI, R. A. (2009). Bioinformatics enrichment tools: Paths toward the comprehensive functional analysis of large gene lists. *Nucleic Acids Res.* **37** 1–13.
- JIMÉNEZ, R. and YUKICH, J. E. (2002). Asymptotics for statistical distances based on Voronoi tessellations. *J. Theoret. Probab.* **15** 503–541. [MR1898817](#)
- JIN, J. and CAI, T. T. (2007). Estimating the null and the proportional of nonnull effects in large-scale multiple comparisons. *J. Amer. Statist. Assoc.* **102** 495–506. [MR2325113](#)
- KIM, J. H., DHANASEKARAN, S. M., MEHRA, R., TOMLINS, S. A., GU, W., YU, J., KUMAR-SINHA, C., CAO, X., DASH, A., WANG, L. et al. (2007). Integrative analysis of genomic aberrations associated with prostate cancer progression. *Cancer Res.* **67** 8229–8239.
- LOUGHIN, T. M. (2004). A systematic comparison of methods for combining p -values from independent tests. *Comput. Statist. Data Anal.* **47** 467–485. [MR2086483](#)
- MURALIDHARAN, O. (2010). An empirical Bayes mixture method for effect size and false discovery rate estimation. *Ann. Appl. Stat.* **4** 422–438. [MR2758178](#)
- OKABE, A., BOOTS, B., SUGIHARA, K. and CHIU, S. N. (2010). *Spatial Tessellations: Concepts and Applications of Voronoi Diagrams*, 2nd ed. Wiley, Chichester.
- OLIVA, A., ROSEBROCK, A., FERREZUELO, F., PYNE, S., CHEN, H., SKIENA, S., FUTCHER, B. and LEATHERWOOD, J. (2005). The cell cycle-regulated genes of *Schizosaccharomyces pombe*. *PLoS Biol.* **3** e225.
- OWEN, A. B. (2009). Karl Pearson’s meta-analysis revisited. *Ann. Statist.* **37** 3867–3892.
- PHILLIPS, D. and GHOSH, D. (2014a). Supplement to “Testing the disjunction hypothesis using Voronoi diagrams with applications to genetics.” DOI:[10.1214/13-AOAS707SUPPA](#).

- PHILLIPS, D. and GHOSH, D. (2014b). Supplement to “Testing the disjunction hypothesis using Voronoi diagrams with applications to genetics.” DOI:10.1214/13-AOAS707SUPPB.
- POLLACK, J. R., SØRLIE, T., PEROU, C. M., REES, C. A., JEFFREY, S. S., LONNING, P. E., TIBSHIRANI, R., BOTSTEIN, D., BØRRESEN-DALE, A.-L. and BROWN, P. O. (2002). Microarray analysis reveals a major direct role of DNA copy number alteration in the transcriptional program of human breast tumors. *Proc. Natl. Acad. Sci. USA* **99** 12963–12968.
- POUNDS, S. B. (2006). Estimation and control of multiple testing error rates for microarray studies. *Brief. Bioinf.* **7** 25–36.
- POUNDS, S. and MORRIS, S. W. (2003). Estimating the occurrence of false positives and false negatives in microarray studies by approximating and partitioning the empirical distribution of p -values. *Bioinformatics* **19** 1236–1242.
- PYKE, R. (1965). Spacings. *J. R. Stat. Soc. Ser. B Stat. Methodol.* **27** 395–449.
- STOREY, J. D. (2002). A direct approach to false discovery rates. *J. R. Stat. Soc. Ser. B Stat. Methodol.* **64** 479–498. MR1924302
- STOUFFER, S. A., SUCHMAN, E. A., DEVINNEY, L. C., STAR, S. A. and WILLIAMS JR., R. M. (1949). *The American Soldier: Adjustment During Army Life. Stud. Soc. Psychol. World War II* **1**. Princeton Univ. Press, Princeton.
- STRIMMER, K. (2008). A unified approach to false discovery rate estimation. *BMC Bioinformatics* **9** 303.
- TAGNON, H. J., WHITMORE, W. F. and SHULMAN, N. R. (1952). Fibrinolysis in metastatic cancer of the prostate. *Cancer* **5** 9–12.
- TONON, G., WONG, K.-K., MAULIK, G., BRENNAN, C., FENG, B., ZHANG, Y., KHATRY, D. B., PROTOPOPOV, A., YOU, M. J., AGUIRRE, A. J. et al. (2005). High-resolution genomic profiles of human lung cancer. *Proc. Natl. Acad. Sci. USA* **102** 9625–9630.
- TSAFRIR, D., BACOLOD, M., SELVANAYAGAM, Z., TSAFRIR, I., SHIA, J., ZENG, Z., LIU, H., KRIER, C., STENGEL, R. F., BARANY, F. et al. (2006). Relationship of gene expression and chromosomal abnormalities in colorectal cancer. *Cancer Res.* **66** 2129–2137.
- TURNER, R. (2013). deldir: Delaunay triangulation and Dirichlet (Voronoi) tessellation.
- WILKINSON, B. (1951). A statistical consideration in psychological research. *Psychol. Bull.* **48** 156.
- ZACHARSKI, L. R., WOJTUKIEWICZ, M. Z., COSTANTINI, V., ORNSTEIN, D. L., MEMOLI, V. A. et al. (1992). Pathways of coagulation/fibrinolysis activation in malignancy. *Semin. Thromb. Hemost.* **18** 104.

DEPARTMENT OF STATISTICS
PENNSYLVANIA STATE UNIVERSITY
UNIVERSITY PARK, PENNSYLVANIA 16802
USA
E-MAIL: dlp245@psu.edu
ghoshd@psu.edu