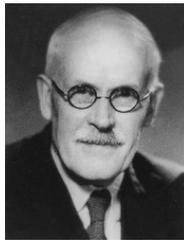


True and false discoveries with independent and sequential e-values

Vladimir Vovk

Ruodu Wang



Users of these tests speak of the 5 per cent. point [p-value of 5%] in much the same way as I should speak of the $K = 10^{-1/2}$ point [e-value of $10^{1/2}$], and of the 1 per cent. point [p-value of 1%] as I should speak of the $K = 10^{-1}$ point [e-value of 10].

Project “Hypothesis testing with e-values”

Working Paper #4

First posted March 1, 2020. Last revised August 13, 2024.

Project web site:
<http://alrw.net/e>

Abstract

In this paper we use e-values in the context of multiple hypothesis testing assuming that the base tests produce independent, or sequential, e-values. Our simulation and empirical studies and theoretical considerations suggest that, under this assumption, our new algorithms are superior to the known algorithms using independent p-values and to our recent algorithms designed for e-values without the assumption of independence.

Contents

1	Introduction	1
2	E-values for hypotheses testing and confidence estimation	2
3	Multiple testing of a single null hypothesis	3
4	Discovery matrices for independent and sequential e-values	4
5	A simulation study with independent e-values	10
6	An empirical study with independent e-values	14
7	Towards efficient algorithms for U_2 and other U_n	17
8	When can we expect the U_2 merging function to be effective?	19
9	Conclusion	21
	Acknowledgments	21
	References	21

1 Introduction

The notion of a p-value has been widely criticized recently as the basis of statistical hypothesis testing (see, e.g., [23]). Bayes factors provide an alternative approach (see, e.g., [1, Sect. 2]). In this paper we continue the study of e-values, a promising new tool for testing, in addition to p-values and Bayes factors. For recent work on hypothesis testing with e-values, as well as their connections and differences to methods based on p-values, see, e.g., [19, 15, 22, 12].

Two important mathematical advantages of e-values over p-values are that the arithmetic average of several e-values is always an e-value and that the product of several independent e-values is an e-value. The second property generalizes to the case of sequential e-values, which is a far-reaching generalisation of independent e-values to be discussed later in the paper. Both properties are useful in multiple hypothesis testing.

Our previous papers [19, 20] applying e-values to multiple hypothesis testing did not make any assumptions about the base e-values and relied on the arithmetic average of e-values being an e-value. Using arithmetic averaging is very natural in this case since the arithmetic mean essentially dominates any symmetric function for merging e-values [19, Proposition 3.1].

In this paper we assume that the base e-values are independent or at least sequential. This gives us extra freedom in combining e-values, which we use for constructing confidence bounds for the number of true discoveries in multiple hypothesis testing. These confidence bounds are summarised in the form of “discovery matrices”, which were introduced in [20] in the case of arbitrary dependence between the base e-values. In this paper we design similar procedures for sequential e-values and demonstrate the power of our procedures in simulation and empirical studies. The assumption that the base e-values are sequential (or independent) is, of course, a downside of the procedures developed in this paper; however, under those assumptions, the gain in power is very significant.

We start the main part of the paper in Sect. 2 from basic definitions related to e-values and confidence bounds. Multiple hypothesis testing is usually understood as testing multiple hypotheses, and this is our main object of interest in this paper. However, Sect. 3 deals with repeated testing of a single hypothesis and introduces functions for merging e-values, which is a first step towards testing multiple hypotheses. In Sect. 4 we introduce discovery matrices. Section 5 is devoted to a simulation study, and Sect. 6 to an empirical study of discovery matrices. In Sect. 7 we develop a computationally more efficient version of our main procedure for computing discovery matrices. Statistical properties of this procedure are briefly discussed in Sect. 8; this is just a first attempt of theoretical analysis. Section 9 concludes and lists some directions of further research.

2 E-values for hypotheses testing and confidence estimation

To make our exposition self-contained, we start from basic definitions. For further information on e-values, see our previous papers [19, 20].

In this paper we fix an underlying probability space $(\Omega, \mathcal{A}, \mathbb{P})$ (which we rarely mention explicitly). We will use the notation $\mathbb{E}(f)$ for the expectation $\int f d\mathbb{P}$ of an extended random variable $f : \Omega \rightarrow [-\infty, \infty]$ w.r. to the true data-generating distribution \mathbb{P} . More generally, we will write $\mathbb{E}_Q(f)$ for the expectation $\int f dQ$ of f w.r. to an arbitrary probability measure Q on (Ω, \mathcal{A}) . The expectation $\mathbb{E}_Q(f)$ always exists (but may be equal to ∞) when f is nonnegative, that is $f : \Omega \rightarrow [0, \infty]$.

An *e-variable* w.r. to a probability measure Q on (Ω, \mathcal{A}) is a nonnegative extended random variable $E : \Omega \rightarrow [0, \infty]$ such that $\mathbb{E}_Q(E) \leq 1$. A large value of $E(\omega)$ for the realized outcome ω is interpreted as casting doubt on Q being the true data-generating distribution \mathbb{P} . Indeed, by Markov's inequality, $Q(E \geq c) \leq 1/c$ for any $c > 1$, and so E can take large values only with a small Q -probability. This interpretation assumes, of course, that E was chosen in advance of the statistical experiment. An *e-value* is the value $E(\omega)$ taken by the e-variable.

A more standard way of testing statistical hypotheses is to use *p-values*, defined to be the values taken by *p-variables*, i.e., nonnegative random variables P such that, for any $\alpha \in (0, 1)$, $Q(P \leq \alpha) \leq \alpha$. There are numerous ways of converting p-values to e-values and a natural way of converting e-values to p-values (see, e.g., [19, Sect. 2]), but in order to compare the strength of evidence against the null hypothesis provided by e-values and p-values we will use Jeffrey's rules of thumb that we will describe in Sect. 5.

Alongside the true data-generating probability measure \mathbb{P} on the underlying measurable space (Ω, \mathcal{A}) we will consider other probability measures and will use the notation \mathcal{Q} for the set of all probability measures on (Ω, \mathcal{A}) . An *e-test* is a family of e-variables $(E_Q \mid Q \in \mathcal{Q})$, E_Q being an e-variable w.r. to Q . Assuming the e-test E had been chosen before the experiment, the interpretation of $E_Q(\omega)$ is that it measures lack of agreement between an outcome ω and a putative explanation Q . Namely, suppose $E_Q(\omega)$ is large; then the disagreement between Q and ω has two sides:

- if ω happens, we do not regard Q to be feasible as a data-generating distribution;
- if we believe that Q is the data-generating distribution ($Q = \mathbb{P}$), we do not expect ω to happen.

If $E_Q(\omega) \geq \alpha$, we will say that ω is α -*strange* under Q . Otherwise, ω and Q *agree* at level α .

2.1 Confidence regions

We are often interested not in the data-generating distribution \mathbb{P} itself but in the value $g(\mathbb{P})$ of a function g on it; for example g can map the probability measures on the real line \mathbb{R} to their medians. In confidence estimation, $g(\mathbb{P})$ is usually interpreted as the value of a parameter corresponding to \mathbb{P} ; we will refer to g as a *parameter function*. For a given e-test E , parameter function $g : \mathcal{Q} \rightarrow \Theta$, *significance level* $\alpha > 0$, and realized outcome $\omega \in \Omega$, the corresponding *confidence region* is defined as

$$\Gamma_\alpha^g(\omega) := \{g(Q) \mid E_Q(\omega) < \alpha\}. \quad (1)$$

A typical value of α used in this paper is 10; in this case $\Gamma_\alpha^g \subseteq \Theta$ (Θ being the chosen *parameter space*) is the set of all parameter values that agree with the realized outcome at level 10.

Notice that the confidence regions (1) are valid simultaneously for all functions $g : \mathcal{Q} \rightarrow \Theta$ (let us assume that the parameter space Θ is fixed, although our statement of simultaneous validity is also true for variable Θ), provided we are using the same e-test for all g . Namely, $g(Q) \in \Gamma_\alpha^g(\omega)$ for all g unless ω is α -strange under Q .

2.2 Necessity and possibility measures

These are different ways of packaging the same information as that available in confidence regions (1). Given an e-test E , we define the *necessity measure*

$$\square^g(B \mid \omega) := \inf_{Q: g(Q) \notin B} E_Q(\omega)$$

and the *possibility measure*

$$\diamond^g(B \mid \omega) := \inf_{Q: g(Q) \in B} E_Q(\omega) = \square^g(B^c \mid \omega), \quad (2)$$

where $B \subseteq \Theta$ is a set of parameter values and $B^c := \Theta \setminus B$ is its complement.

The intuition behind $\square^g(B \mid \omega)$ is given in terms of a Fisher-type disjunction: $g(\mathbb{P}) \in B$ unless ω is $\square^g(B \mid \omega)$ -strange under \mathbb{P} (cf. [7, Sect. III.1]). For example, we expect $g(\mathbb{P}) \in B$ if $\square^g(B \mid \omega)$ is large. This also gives an interpretation of $\diamond^g(B \mid \omega)$ in view of the duality

$$\diamond^g(B \mid \omega) = \square^g(B^c \mid \omega).$$

Namely, $g(\mathbb{P}) \in B$ is impossible unless ω is $\diamond^g(B \mid \omega)$ -strange. (So we could also call (2) an impossibility measure.) For example, we expect $g(\mathbb{P}) \notin B$ if $\diamond^g(B \mid \omega)$ is large.

3 Multiple testing of a single null hypothesis

A measurable function $F : [0, \infty)^K \rightarrow [0, \infty)$ for an integer $K \geq 1$ is an *ie-merging function* if, for any probability space and any independent e-variables

E_1, \dots, E_K on it, the extended random variable $F(E_1, \dots, E_K)$ is an e-variable (when it exists, i.e., when E_1, \dots, E_K take finite values).

An important relaxation of independence of e-variables E_1, \dots, E_K is the requirement that they be *sequential*: for each $k \in \{1, \dots, K\}$,

$$\mathbb{E}(E_k \mid E_1, \dots, E_{k-1}) \leq 1.$$

Sequential e-variables typically describe a sequential process of testing the null hypothesis \mathbb{P} : E_k should be a valid e-variable even when E_1, \dots, E_{k-1} are known. A measurable function $F : [0, \infty)^K \rightarrow [0, \infty)$ is an *se-merging function* if, for any probability space and any sequential e-variables E_1, \dots, E_K on it, $F(E_1, \dots, E_K)$ is an e-variable. Of course, all se-merging functions are ie-merging functions, but the converse is not true [21, Example 2].

We extend these merging functions to be of the form $F : [0, \infty]^K \rightarrow [0, \infty]$ in a canonical way (described in [19], Proposition C.1 in the Online Supplement): if any of the arguments of F is ∞ , set F to ∞ .

Important examples of se-merging functions [19] are

$$U_n(e_1, \dots, e_K) := \frac{1}{\binom{K}{n}} \sum_{\{k_1, \dots, k_n\} \subseteq \{1, \dots, K\}} e_{k_1} \dots e_{k_n}, \quad n \in \{1, \dots, K\}. \quad (3)$$

We will refer to them as the *U-statistics* (they are the standard U-statistics with product as kernel). The statistics U_1 play a special role since they belong to the narrower class of *e-merging functions*, meaning that $U_1(E_1, \dots, E_k)$ is an e-variable whenever E_1, \dots, E_k are e-variables (not necessarily sequential). Multiple hypothesis testing using U_1 was explored in [19, 20], and in this paper we will often be interested in U_2 .

An assumption that we will need to make about merging functions in this paper is that they are increasing in each argument and are symmetric (more fully, permutation-symmetric, i.e., not depending on the order of their arguments). In particular, Algorithm 1 below will assume that its underlying merging function F is increasing and symmetric. While this assumption appears to be very natural for ie-merging functions, it becomes more restrictive for se-merging functions: see Example 1 in [21].

4 Discovery matrices for independent and sequential e-values

In this section we consider the problem of testing K statistical hypotheses H_k , $k = 1, \dots, K$, for some $K \in \{2, 3, \dots\}$. Each hypothesis H_k may be *simple*, $H_k \in \mathcal{Q}$, or *composite*, $H_k \subseteq \mathcal{Q}$. We assume, without loss of generality, that H_k is composite (a simple hypotheses $Q \in \mathcal{Q}$ can be interpreted as the composite hypothesis $\{Q\}$).

For each k we are testing the null hypothesis H_k using an e-variable E_k . Namely, we are given a sequence E_1, \dots, E_K of extended random variables such

that each E_k , $k = 1, \dots, K$, is an e-variable for testing H_k in the sense of satisfying $\mathbb{E}_Q(E_k) \leq 1$ for all $Q \in H_k$. Therefore, E_k is an e-variable w.r. to any $Q \in H_k$, but it does not have to be an e-variable w.r. to the true data-generating measure \mathbb{P} .

4.1 Discovery vectors

We are interested in the set of $k \in \{1, \dots, K\}$ for which $\mathbb{P} \in H_k$, i.e., H_k is a true hypothesis (which also determines the set of indices for which the hypothesis is false). Our first goal is to define, for each set $R \subseteq \{1, \dots, K\}$, a confidence region for the number of $k \in R$ for which H_k is false. Our confidence region will always be of the form $\{L, \dots, K\}$ for some lower bound L (we only have a lower confidence bound since $\mathbb{E}(E_k) \leq 1$ may be true even if H_k is not a true hypothesis).

Suppose that H_1, \dots, H_K are scientific hypotheses that an experimental scientist is interested in. After observing the e-values e_1, \dots, e_K , where e_k is the realized value of the e-variable E_k , the scientist comes up with a *rejection set* $R \subseteq \{1, \dots, K\}$ containing the indices of the hypotheses that she decides to reject based on e_1, \dots, e_K . The elements of R are referred to as *discoveries* (i.e., they are discoveries as claimed by the scientist). A discovery $k \in R$ is a *true discovery* if $\mathbb{P} \notin H_k$, and it is a *false discovery* if $\mathbb{P} \in H_k$. The motivation for this terminology is that H_k is regarded to be the default hypothesis saying that an interesting effect does not exist [14]; rejecting the null hypothesis is then a discovery.

A natural way to choose the rejection set R is to include in it a number of k with the largest e_k , but the scientist sometimes might want to include all k connected by some common theme [11, Sect. 4.1]. We will start from discussing arbitrary R and then will move on to the R corresponding to the largest e-values.

For each $Q \in \mathcal{Q}$, let

$$I_Q := \{k \in \{1, \dots, K\} \mid Q \in H_k\}$$

be the indices of the true hypotheses under Q . Given a rejection set R , we are interested in the parameter

$$g_R(Q) := |R \setminus I_Q|,$$

which is the number of true discoveries in R under Q . While in this paper we concentrate on the parameter function g_R , this function can be generalized in various directions; see, e.g., [20, Remark 6.1].

We are interested in lower confidence bounds on the number of true discoveries. For that, in order to apply the recipe (1), we need an e-test. A natural way to obtain an e-test is to apply a sequence of symmetric e-merging functions $F_{K'} : [0, \infty)^{K'} \rightarrow [0, \infty)$, $K' \in \{2, \dots, K\}$, at each $Q \in \mathcal{Q}$:

$$E_Q := F_{|I_Q|}((E_k)_{k \in I_Q}) = F((E_k)_{k \in I_Q}), \quad (4)$$

where $(E_k)_{k \in I_Q}$ stands for the sequence of length $|I_Q|$ whose elements are E_k with $k \in I_Q$ in the ascending order of k (although the order of the elements of this sequence does not matter since $F_{K'}$ are assumed symmetric). We define F_0 and F_1 in (4) (needed when $|I_Q| \in \{0, 1\}$) separately; namely, $F_0 := 1$ and $F_1(e) := e$. The second “=” in (4) introduces our convention of omitting the lower index K' in the notation $F_{K'}$; indeed, the lower index is redundant as it can be recovered from the number of arguments of $F_{K'} = F_{K'}(e_1, \dots, e_{K'})$.

In order to obtain tighter confidence regions, we assume that, for any probability measure $Q \in \mathcal{Q}$, the random e-variables E_k , $k \in I_Q$, are independent under Q . It is clear that in this case (4) is an e-test provided F is a symmetric ie-merging function. This will be the main case considered in this paper.

We gave examples of useful ie-merging functions in Sect. 3. However, for the use in (4) we need to generalize (3) to the case $n > K \geq 1$ (remember that we have $K' := |I_Q|$ in place of K when using (3) in the context of (4), and so K' can be a small number even for a large number of the null hypotheses). Let us set

$$U_n(e_1, \dots, e_K) := U_K(e_1, \dots, e_K) = e_1 \dots e_K, \quad n > K. \quad (5)$$

This convention is not needed in the case $n = 2$, which is our main object of study, since it was already made implicitly when defining F_1 .

For each $j \in \{0, \dots, K - 1\}$, we are interested in the possibility measure

$$\diamond^{gR}(\{0, \dots, j\} \mid \omega),$$

where $\omega \in \Omega$ is the realized outcome, which we usually omit, as is customary in probability theory. The interpretation of $\diamond^{gR}(\{0, \dots, j\})$ is that the number of true discoveries exceeds j unless the realized outcome is $\diamond^{gR}(\{0, \dots, j\})$ -strange.

A more explicit representation of $\diamond^{gR}(\{0, \dots, j\})$ is:

$$\begin{aligned} \diamond^{gR}(\{0, \dots, j\}) &= \inf_{Q \in \mathcal{Q}: |R \setminus I_Q| \leq j} E_Q = \min_{Q \in \mathcal{Q}: |R \setminus I_Q| \leq j} F((E_k)_{k \in I_Q}) \\ &\geq \min_{I \subseteq \{1, \dots, K\}: |R \setminus I| \leq j} F((E_i)_{i \in I}) =: D_F^R(j). \end{aligned} \quad (6)$$

The inequality in (6) may be strict (when some subsets of $\{1, \dots, K\}$ cannot be represented in the form I_Q for any $Q \in \mathcal{Q}$), but in interesting cases we have an equality there. The symbol “=” means that $D_F^R(j)$ is being defined.

The vector $(D_F^R(j))_{j=0}^{K-1}$ is the *discovery vector* for R (in [19] we referred to it as regularized discovery e-vector). We will now introduce “discovery matrices”, whose rows are discovery vectors.

4.2 Discovery matrices

In this subsection we impose restrictions on the ie-merging function F in (6) which will allow us to restrict our attention to a relatively small number of rejection sets. Namely, as mentioned at the end of Sect. 3, we assume that F is increasing in each of its arguments and symmetric. Suppose the e-values are listed in the decreasing order, $e_1 \geq \dots \geq e_K$. Then the rejection sets

$$R_r := \{1, \dots, r\}, \quad r \in \{1, \dots, K\},$$

Algorithm 1 Discovery matrix (lower triangular)

Input: ie-merging (or se-merging) functions F_k , $k \in \{1, \dots, K\}$.

Input: a decreasing sequence of e-values $e_1 \geq \dots \geq e_K$.

```

1: for  $r = 1, \dots, K$ 
2:   for  $j = 0, \dots, r - 1$ 
3:      $S_{r,j} := \{j + 1, \dots, r\}$ 
4:      $DM_{r,j}^F := F((e_i)_{i \in S_{r,j}})$ 
5:     for  $k = r + 1, \dots, K$ 
6:        $e := F((e_i)_{i \in S_{r,j} \cup \{k, \dots, K\}})$ 
7:       if  $e < DM_{r,j}^F$ 
8:          $DM_{r,j}^F := e$ 

```

form a complete family of rejection sets: for any r and any rejection set R of size r we have, for all $j \in \{0, \dots, r - 1\}$, $D_{F^r}^{R_r}(j) \geq D_F^R(j)$.

The *discovery matrix* is a lower triangular matrix with the entries

$$DM_{r,j} = DM_{r,j}^F := D_{F^r}^{R_r}(j), \quad r = 1, \dots, K, \quad j = 0, \dots, r - 1;$$

we often drop the upper index F in $DM_{r,j}^F$ when it is clear from the context. Algorithm 1 is one way of constructing a discovery matrix, under an additional assumption (cf. (8) below), based on a family of ie-merging functions F_k , $k \in \{2, \dots, K\}$. The algorithm implements the equality between the extreme terms of

$$DM_{r,j}^F = \min_{I \subseteq \{1, \dots, K\}: |R^r \setminus I| \leq j} F((E_i)_{i \in I}) = \min_{I \subseteq \{1, \dots, K\}: |R^r \setminus I| = j} F((E_i)_{i \in I}). \quad (7)$$

The first equality in (7) follows from the last equality (“=”) in (6), and the second equality in (7) follows from the following natural condition of consistency between different merging functions in the family (F_k) : if $\mathbf{e} \in [0, \infty)^k$ for $k \in \{1, \dots, K - 1\}$ and $e \in [0, \infty)$, then

$$e \geq \max(\mathbf{e}) \implies F(\mathbf{e}, e) \geq F(\mathbf{e}). \quad (8)$$

This is the condition under which Algorithm 1 is valid.

Remark 4.1. Algorithm 1 is a version of Algorithm 2 in [20]. In both algorithms the e-values are assumed to be ordered (without loss of generality, since we assume F to be symmetric), but while in Algorithm 2 in [20] the order is ascending, in Algorithm 1 it is descending. Another difference is that in this paper (as in the arXiv version of [20]) we use the version of the discovery matrix described in [20, Remark 6.3]. Finally, in [20] we were mainly interested in the arithmetic-mean merging function (because it essentially dominates all symmetric e-merging functions with no dependence assumption), while in this paper we are interested in a wider range of merging functions by considering independent or sequential e-variables.

The validity of Algorithm 1 follows from the fact that the min in

$$\min_{I \subseteq \{1, \dots, K\}: |R \setminus I| = j} F((E_i)_{i \in I})$$

(cf. (7)) is attained at I of the form

$$S_{r,j} \cup \{k, \dots, K\} = \{j+1, \dots, r\} \cup \{k, \dots, K\}.$$

This follows immediately from the ie-merging function F being increasing in all arguments and symmetric.

Discovery matrices satisfy useful properties of monotonicity given in the following proposition (which is proved as part of the proof of Proposition 6.3 in the arXiv version of [20]). The following proposition assumes the standard way of presenting matrices (as in Figures 1–5 below).

Proposition 4.2. *Suppose the family (F_k) of ie-merging functions satisfies (8). Then $DM_{r,j}$ is*

- *decreasing along the rows: $DM_{r,j'} \leq DM_{r,j}$ when $j' > j$;*
- *increasing down the columns: $DM_{r',j} \geq DM_{r,j}$ when $r' > r$;*
- *decreasing in the Southeast direction: $DM_{r+i,j+i} \leq DM_{r,j}$ when $i > 0$.*

(Even if (8) is violated, these properties are satisfied for the regularized version (9) defined below.)

Condition (8), despite looking very natural, is not satisfied, strictly speaking, even for the U_2 ie-merging function. For example, taking $\mathbf{e} = (e_1)$ of length 1, we have

$$U_2(e_1, e) = e_1 e < e_1 = U_2(e_1) = U_1(e_1)$$

(remember our convention (5)) for some e_1 and e such that $e > e_1$: it suffices to take $e < 1$. In this case, however, both e_1 and $e_1 e$ are useless as e-values, and in practice we can regard (8) as satisfied for U_2 (it is natural to ignore entries below 1 in discovery matrices, or to replace them by 1).

If (8) is violated, we should regularize the output of Algorithm 1 by redefining

$$DM_{r,j}^F := \min(DM_{r,0}^F, \dots, DM_{r,j}^F). \quad (9)$$

Then the equality between the extreme terms of (7) will always hold, and so $DM_{r,j}$ will be a lower bound on $\diamond^{gR}(\{0, \dots, j\})$.

The computation time of Algorithm 1 depends on the computation time of the family of ie-merging function F_k , $k = 1, \dots, K$. In the cases of primary interest to us, namely U_n for a fixed n , the computation time for each of F_k is $O(K)$ (and even $O(k)$). For U_1 this is obvious, and for U_2 it suffices to represent it in the form

$$U_2(e_1, \dots, e_k) = \frac{1}{K(K-1)} ((e_1 + \dots + e_k)^2 - (e_1^2 + \dots + e_k^2)).$$

For a general fixed n , use induction in n .

Assuming that F_k , $k = 1, \dots, K$, are computed in linear time, $O(K)$, each row of the discovery matrix can be computed by Algorithm 1 in time $O(K^3)$: computing each e-value (line 8) takes time $O(K)$, each inner loop (line 5) takes time $O(K^2)$, and so each middle loop (line 2) takes time $O(K^3)$. Therefore, computing the whole discovery matrix takes time $O(K^4)$. Applying regularization (9) on top of these calculations, if needed, only takes time $O(K)$, and so does not change the overall time ($O(K^3)$ in the case of a row of the discovery matrix or $O(K^4)$ in the case of the whole discovery matrix).

4.3 Beyond independent e-values

So far we have discussed the case of independent e-variables E_k (more precisely, we have assumed the Q -independence of $(E_k \mid Q \in H_k)$ under any $Q \in \mathcal{Q}$). Replacing ie-merging functions by se-merging functions, we can assume, instead, that E_1, \dots, E_K are sequential, in the sense of satisfying $\mathbb{E}_Q(E_k \mid E_1, \dots, E_{k-1}) \leq 1$ whenever $Q \in H_k$. To check that (4) is indeed an e-test, apply the tower property of conditional expectations: the latter implies that $(E_k)_{k \in I_Q}$ are sequential e-variables under Q whenever E_1, \dots, E_K are.

In Sections 5 and 6 we will apply Algorithm 1 to K independent e-values. A typical scenario in which independent e-values arise in K statistical experiments is where there are reasons to believe that there is no or little interference or sample sharing between different experiments. The experiments can be conducted at the same time or at different times.

Remark 4.3. It is interesting that, in the case of U_2 considered in the following two sections, the assumption of independence between E_1, \dots, E_K can be relaxed in a way different from this paper's assumption that E_1, \dots, E_K are sequential. Instead, we can assume that the covariances $\text{cov}(E_i, E_j)$, $i \neq j$, are all nonpositive. Indeed, in this case

$$\begin{aligned} \mathbb{E}(E_i E_j) &= \mathbb{E}((\mathbb{E}E_i + (E_i - \mathbb{E}E_i))(\mathbb{E}E_j + (E_j - \mathbb{E}E_j))) \\ &= (\mathbb{E}E_i)(\mathbb{E}E_j) + \text{cov}(E_i, E_j) \leq 1. \end{aligned}$$

Chi et al. [3, Sect. 2] describe interesting practical situations where this assumption is realistic. This includes sampling without replacement. Suppose there is a very large fixed population, and K scientists each take a different sample from the population without replacement. Then any increasing statistics (e.g., sample means) of these K subsamples are negatively associated, and their correlation is nonpositive.

A possible scenario in which sequential (but not independent) e-variables E_1, \dots, E_K arise is where the experiments are performed sequentially, the experiment E_k for $k \in \{2, \dots, K\}$ is started after the experiment E_{k-1} is finished, and the experiment E_k is designed using the results of the previous experiments E_1, \dots, E_{k-1} . It is essential that E_k should be an e-variable conditionally on

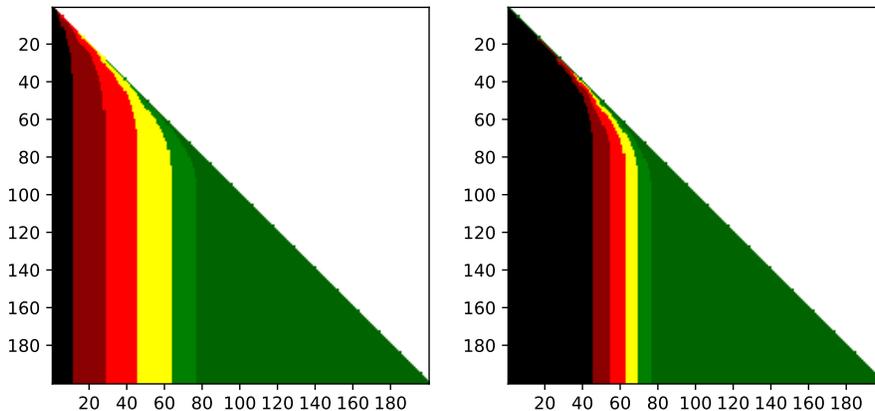


Figure 1: Left panel: the discovery matrix for the U_1 statistic (i.e., arithmetic mean) for 100 false and 100 true null hypotheses. Right panel: the U_2 analogue.

the previous experiments. The computational experiments reported in this paper do not cover this case, but it is important for us mathematically as all our methods are applicable to sequential e-variables.

5 A simulation study with independent e-values

In this section we run Algorithm 1 applied to U_2 and, for comparison, U_1 . As discussed in the previous section, the U_n discovery matrix can be computed in time $O(K^4)$. For $n = 1$, the time can be improved from $O(K^4)$ to $O(K^2)$ [20, Sect. 9]. For $n = 2$, we can improve the time $O(K^4)$ to $O(K^3)$, as we show in Sect. 7 (Algorithm 2), and this is sufficient to cope not only with the case $K = 200$ that we use in our simulation studies in this section but also with K of a few thousand (as used in our empirical studies in the next section, where we compute only part of the discovery matrix).

We generate the base e-values as in [20, Sect. 7]: the null hypotheses are $\mathcal{N}(0, 1)$, $K = 200$, the first 100 observations x are generated from $\mathcal{N}(-3, 1)$, the last 100 from $\mathcal{N}(0, 1)$, all independently, and the base e-variables are the likelihood ratios

$$E(x) := \frac{\exp(-(x+3)^2/2)}{\exp(-x^2/2)} = \exp(-3x - 9/2). \quad (10)$$

The results are shown in Figure 1 (whose left panel is identical to the upper left panel of Figure 4 in [20]); they are much better for U_2 . Our chosen colour scheme, to be described momentarily, entails that “better” means “darker” here.

Each panel shows the lower triangular matrix $DM_{r,j}^F$, the left for $F = U_1$ and the right for $F = U_2$. The colour scheme used in this figure is inspired by

Jeffreys's [13, Appendix B] (as in [20]):

- The entries with $DM_{r,j}^F$ below 1 are shown in dark green; there is no evidence that there are more than j true discoveries among the r hypotheses with the largest e-values.
- The entries $DM_{r,j}^F \in (1, \sqrt{10}) \approx (1, 3.16)$ are shown in light green. For them the evidence is poor.
- The entries $DM_{r,j}^F \in (\sqrt{10}, 10) \approx (3.16, 10)$ are shown in yellow. The evidence is substantial.
- The entries $DM_{r,j}^F \in (10, 10^{3/2}) \approx (10, 31.6)$ are shown in light red. The evidence is strong.
- The entries $DM_{r,j}^F \in (10^{3/2}, 100) \approx (31.6, 100)$ are shown in dark red. The evidence is very strong.
- Finally, the entries $DM_{r,j}^F > 100$ are shown in black, and for them the evidence is decisive.

The interpretation of the two discovery matrices in Figure 1 in terms of confidence regions is particularly convenient: for each row r of a discovery matrix,

- the red (both dark and light), yellow, and green (both dark and light) entries in row r coincide with the confidence region at significance level 100 for the number of true discoveries among the r largest e-values,
- the light red, yellow, and green entries in row r coincide with the confidence region at significance level $10^{3/2}$ for the number of true discoveries among the r largest e-values,
- the yellow and green entries in row r coincide with the confidence region at significance level 10 for the number of true discoveries among the r largest e-values,
- the green entries in row r coincide with the confidence region at significance level $\sqrt{10}$ for the number of true discoveries among the r largest e-values (this is the final interesting case).

A darker discovery matrix (such as the right-hand panel of Figure 1 as compared with its left-hand panel) is preferable since it means tighter confidence regions.

Comparison with methods based on p-values

It is interesting to compare the right panel of Figure 1 to a similar figure for p-values obtained by standard methods. A rigorous way of doing it would be to convert, or as we usually say, *calibrate*, e-values to p-values and vice versa. As we mentioned in Sect. 2, there are numerous ways of calibrating p-values to e-values. However, the only admissible way of calibrating an e-value e to

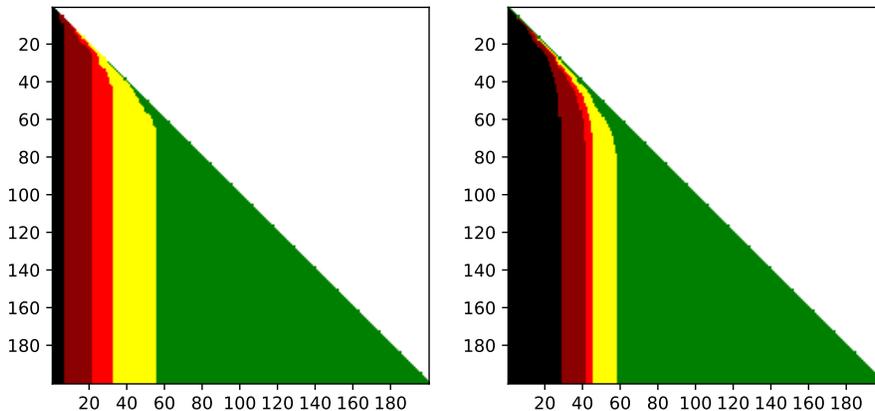


Figure 2: Left panel: the discovery p-matrix for the GWGS procedure with independent p-values. Right panel: the U_2 discovery matrix e-to-p calibrated via $p := \min(1/e, 1)$.

a p-value p is $p := \min(1/e, 1)$. Its validity follows from Markov’s inequality ($\mathbb{E}(1/E \leq \alpha) \leq \alpha$ for any e-variable E), and its domination of any other e-to-p calibrator is stated in [19, Proposition 2.2]. Comparing e-values and p-values is discussed in detail in [20, Sect. 3], where we demonstrate very low “round-trip efficiency” of converting e-values to p-values and back: namely, we start from a highly statistically significant p-value of 0.5%, transform it to an e-value using a popular method [1, (1)], and then transform it back to a p-value by inverting the e-value; the resulting p-value (7.2%) is not even statistically significant. Therefore, in [19] we emphasize informal comparisons.

The left panel of Figure 2 shows the p-values produced by the GWGS procedure, which is designed specifically for p-values. The procedure is very standard [8, 11, 10], but the abbreviation, after the four authors of this procedure, that we use for it was introduced in our previous paper [20]. As presented in that paper, the procedure produces a discovery matrix based on p-values, or as we will say, *discovery p-matrix*; namely, the discovery p-matrix in Figure 2 has the entries

$$\text{DM}_{r,j}^{\text{Simes}} := \max_{I \subseteq \{1, \dots, K\}: |R^r \setminus I| \leq j} F((P_i)_{i \in I}), \quad r = 1, \dots, K, \quad j = 0, \dots, r - 1 \quad (11)$$

(cf. (7)), where F is Simes’s [16] function for merging independent p-values. The underlying p-value P_k for testing the null hypothesis H_k for each of the 200 observations is the optimal one coming from the Neyman–Pearson lemma (and so it is based on (10) as test statistic).

Figure 2 uses what we called Fisher’s scale in [20], but now we extend it by two further thresholds, one of which is 0.5%, as advocated by Benjamin et al.

[2]. Therefore, our colour scheme is:

- The p-values above 5% are shown in green; they are not statistically significant.
- The p-values between 1% and 5% are shown in yellow; they are statistically significant but not highly significant.
- The p-values between 0.5% and 1% are shown in red; they are highly significant (but fail to attain the more stringent criterion of significance advocated in [2]).
- The p-values between 0.1% and 0.5% are shown in dark red.
- The p-values below 0.1% are shown in black; they can be regarded as providing decisive evidence against the null hypothesis (to use Jeffreys's expression in a slightly different context).

Jeffreys [13, Appendix B] gives a crude but convenient rule of thumb for comparing Bayes factors and p-values. E-values are closely related to Bayes factors, and in fact e-values are Bayes factors (and vice versa) in the case of simple null hypotheses; see, e.g., [12, Sect. 1.1.3] for a detailed discussion. According to Jeffreys's rule of thumb as applied to e-values and p-values, a p-value of 5% corresponds to an e-value of $10^{1/2} \approx 3.16$, and a p-value of 1% corresponds to an e-value of 10. This and similar informal correspondences suggested by other authors (such as I. J. Good) are described in detail in [20, Sect. 3]. Therefore, according to this interpretation of Jeffreys's rule, the discovery matrices and discovery p-matrices using our colour schemes are somewhat comparable; e.g., the light red cells correspond to similar amounts of evidence against the null hypothesis.

Comparing the left panel of Figure 2 to the right panel of Figure 1, we can see that the results produced using e-values are typically much better leading to tighter confidence regions. Remarkably, even after the crude e-to-p calibration $e \mapsto 1/e$ our method produces p-values that look better than the p-values produced by the GWGS procedure: see the right panel of Figure 2.

A simple informal explanation of the superior performance of our method is that when combining independent e-values evidence may multiply. For instance, the product of two "non-substantial" independent e-values (e.g., around 3) leads to a "substantial" e-value (e.g., around 9). In the GWGS procedure, the standard method of combining independent p-values is that of Simes, which will not produce anything smaller than the smallest input p-value. In contrast, our U_2 function (as well as U_n for $n \geq 3$) is able to produce larger output e-values than the largest individual input. This observation suggests that other versions of the GWGS procedure, for instance combining a U-statistic with Fisher's method [6], might be more powerful than the standard method in certain applications.

6 An empirical study with independent e-values

In this section we will use the `prostate` dataset first described by Singh et al. [17]; it is also analyzed in [4, Chap. 2] and then in [5, Chap. 15]. This dataset represents a 6033×102 matrix whose rows correspond to 6033 genes and columns correspond to 102 men. The first 50 men are healthy controls and the remaining 52 are patients with prostate cancer. Each entry $x_{k,j}$ of the matrix represents the activity of the k th gene in the j th man.

For each gene we are interested in whether its activity is different in the patients and the healthy controls. So we have 6033 null hypotheses of no difference to test. Efron [4, Chap. 2] makes the assumption of independence of his test statistics for testing those null hypotheses (but he also analyses shortfalls of this assumption in [4, Sect. 2.5]), along with several other substantial assumptions, such as the Gaussian distribution of the genes' activities. In this paper, however, we will avoid making any other assumptions apart from independence.

Following [20, Sect. 8] (which, however, considered a different dataset), we compute the base e-values as

$$e_k := \frac{T_k}{\frac{1}{B+1} \left(\sum_{b=1}^B T_k^{(b)} + T_k \right)}, \quad k \in \{1, \dots, 6033\}, \quad (12)$$

where T_k is the *nonconformity score* computed as described in the next paragraph from the k th row of the data matrix, $T_k^{(b)}$ is the nonconformity score computed from the same row with randomly permuted labels (independently for different b), and B is the number of permutations.

We define the nonconformity score as $T_k := |t_k|^d$, where $d > 0$ is a parameter of the algorithm,

$$t_k := \frac{\bar{x}_{k,1} - \bar{x}_{k,0}}{s_k}$$

is the two-sample t-statistic for the k th gene, $\bar{x}_{k,1}$ is the average entry in the k th row for the patients, $\bar{x}_{k,0}$ is the average entry in the k th row for the healthy controls,

$$s_k^2 := \sum_j (x_{k,j} - \bar{x}_{k,y_j})^2$$

is the sample variance of row k (ignoring a constant factor, which cancels out when computing (12)), and y_j is the label (1 for the patients and 0 for the healthy controls).

The left panel of Figure 3 gives the top-left 200×200 corner of the discovery matrix for $d := 10$ and $B := 10000$. The result is much better than for the U_1 discovery matrix: see the left panel of Figure 4. The price to pay for U_2 giving tighter confidence regions than U_1 is, of course, the reliance of the former on the independence assumption for the base e-variables.

In the right panels of Figures 3 and 4 we give analogous plots but with the proper e-values (12) replaced by their commonly used simplified versions

$$e_k := \frac{T_k}{\frac{1}{B} \sum_{b=1}^B T_k^{(b)}}, \quad k \in \{1, \dots, 6033\}. \quad (13)$$

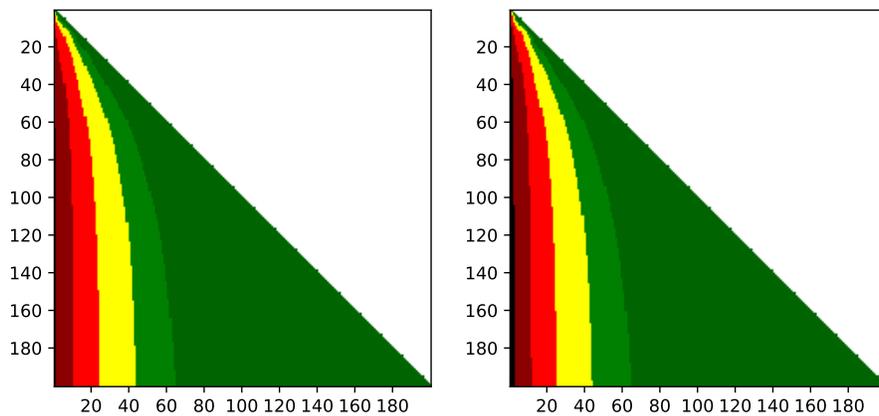


Figure 3: Left panel: the top-left 200×200 corner of the U_2 discovery matrix for the `prostate` dataset for $B := 10000$, using Jeffreys's thresholds. Right panel: its simplified version (based on (13)) that is only approximately valid.

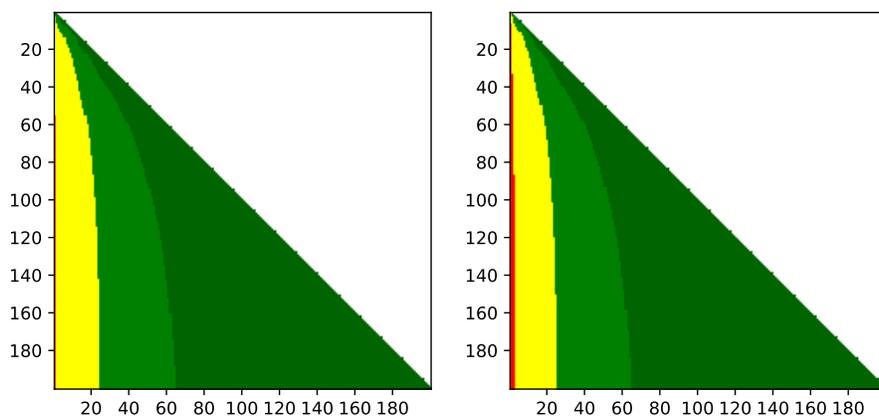


Figure 4: The analogue of Figure 3 for the U_1 discovery matrix.

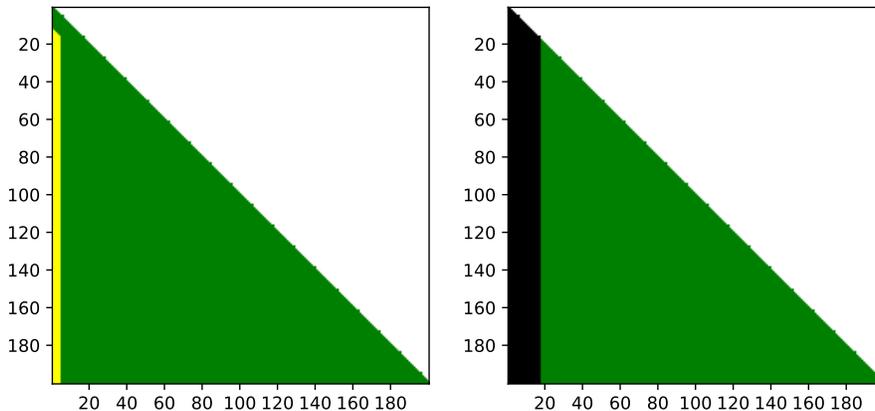


Figure 5: The analogue of Figure 3 for the discovery p-matrix using the GWGS method with independent p-values.

We will call (13) *simplified e-values*, but in fact they are not bona fide e-values. The similarity between the definitions (12) and (13) suggests that simplified e-values are approximately valid, but sometimes their lack of validity is visible. The closeness of the left and right panels in Figures 3 and 4 suggests that the number of iterations $B = 10000$ is sufficiently large.

Comparison with methods based on p-values

Finally, Figure 5 is the analogue of Figures 3 and 4 for the GWGS method (11) applied to the p-values

$$p_k := \frac{\left| \left\{ b \in \{1, \dots, B\} \mid T_k^{(b)} \geq T_k \right\} \right| + 1}{B + 1} \quad (14)$$

(cf. [18, (2.20)]) and their simplified versions

$$p_k := \frac{\left| \left\{ b \in \{1, \dots, B\} \mid T_k^{(b)} \geq T_k \right\} \right|}{B} \quad (15)$$

for $k \in \{1, \dots, 6033\}$. In this case, there is no dependence on d as the p-values only depend on the ranks of T_k . The valid p-values (14) used in the left panel of Figure 5 give a very poor result, and the comparison with the right-hand panel shows that the number of iterations $B = 10000$ is far too small when using p-values: the putative p-values (15) in the right-hand panel are very far from being valid.

Algorithm 2 One row of the discovery matrix $DM = DM^{U_2}$ in time $O(K^2)$

Input: Decreasing sequence of e-values $e_1 \geq \dots \geq e_K$.

Input: Row number $r \in \{1, \dots, K\}$ of the discovery matrix.

```

1:  $s_{K+1} := 0$ 
2: for  $k = K, \dots, 1$ 
3:    $s_k := s_{k+1} + e_k$ 
4: for  $j = r - 1, \dots, 0$ 
5:    $V := V_2(\{j + 1, \dots, r\})$ 
6:    $e := \frac{2V}{(r-j)(r-j-1)}$  (unless  $j = r - 1$ )
7:   for  $k = K, \dots, r + 1$ 
8:      $V' := V_2(\{j + 1, \dots, r\} \cup \{k, \dots, K\})$ 
9:      $e' := \frac{2V'}{(r-j+K-k+1)(r-j+K-k)}$ 
10:    if  $e' < e$ 
11:       $e := e'$ 
12:     $DM_{r,j} := e$ 

```

7 Towards efficient algorithms for U_2 and other U_n

In this appendix we will see that each row of the discovery matrix based on U_2 can be computed in time $O(K^2)$, and so the computation of the full discovery matrix takes time $O(K^3)$. This is not as good as the $O(K^2)$ algorithm for U_1 given as Algorithm 4 in [20], and the existence of such an algorithm for U_2 remains an open problem.

Algorithm 2 starts (in lines 1–3) from defining an array

$$s_k := e_k + \dots + s_K, \quad k \in \{1, \dots, K + 1\},$$

of partial sums, which is not used in the algorithm explicitly, but we will explain how it enables an efficient update of the variables V and V' in lines 5 and 8. In the algorithm we use the notation

$$V_2(I) := \sum_{i, i' \in I: i < i'} e_i e_{i'}, \quad \emptyset \subset I \subseteq \{1, \dots, K\};$$

this expression is the key component of the U-statistic computed from e_i , $i \in I$. When $|I| = 1$, we will treat the value of $V_2(I)$ as undefined.

The first value of V in line 5, $V_2(\{r\})$, is undefined, and we then set $e := e_r$ instead of the formula given in line 6. The next value of V ,

$$V := V_2(\{r - 1, r\}) = e_{r-1} e_r,$$

is computed from scratch in time $O(1)$, and the following values are computed in time $O(1)$ using the previous value: $V = V_2(\{j + 1, \dots, r\})$ is computed from the previous value as

$$V = V + e_{j+1} (e_{j+2} + \dots + e_r) = V + e_{j+1} (s_{j+2} - s_{r+1}).$$

The very first value of V' computed in line 8 (for $j = r - 1$ and $k = K$) can be found in time $O(1)$ from scratch,

$$V' = V_2(\{r, K\}) = e_r e_K.$$

After that the first value of V' in each execution of the loop starting in line 7 (i.e., V' for $j < r - 1$ and $k = K$) can be found in time $O(1)$ from the current value of V using

$$V' := V + e_K (s_{j+1} - s_{r+1}).$$

For the following iterations of this loop the value of V' can be updated in time $O(1)$ using

$$\begin{aligned} V' &:= V' + e_k (e_{j+1} + \dots + e_r + e_{k+1} + \dots + e_K) \\ &= V' + e_k (s_{j+1} - s_{r+1} + s_{k+1}). \end{aligned}$$

Modifications for U_n , $n > 2$

It is easy (but tiresome) to modify Algorithm 2 so that it computes a row of DM^{U_n} for a fixed $n > 2$ in time $O(K^2)$. Let us consider, for simplicity, the case $n = 3$.

Both entries of V_2 in Algorithm 2 should be changed to V_3 , where

$$V_3(I) := \sum_{i, i', i'' \in I: i < i' < i''} e_i e_{i'} e_{i''},$$

line 6 should be changed to

$$e := \frac{6V}{(r-j)(r-j-1)(r-j-2)}$$

(unless $j = r - 1$ or $j = r - 2$), and line 9 should be changed to

$$e' := \frac{6V'}{(r-j+K-k+1)(r-j+K-k)(r-j+K-k-1)}$$

(unless $j = r - 1$ and $k = K$). Lines 1-3 for computing the array

$$s_k := V_1(\{k, \dots, K\}),$$

where

$$V_1(I) := \sum_{i \in I} e_i,$$

should be complemented by computing the array

$$t_k := V_2(\{k, \dots, K\}), \quad k = K, \dots, 1.$$

The array t can be computed in time $O(K)$ starting from $t_K := 0$ and setting, for $k = K - 1, \dots, 1$,

$$t_k := t_{k+1} + e_k s_{k+1}.$$

In line 5, we can compute V in time $O(1)$ given its previous value using the identity (true unless j is very close to r)

$$V_3(\{j+1, \dots, r\}) = V_3(\{j+2, \dots, r\}) + e_{j+1}V_2(\{j+2, \dots, r\})$$

and the identity

$$V_2(\{j+2, \dots, r\}) = t_{j+2} - t_{r+1} - s_{r+1}(s_{j+2} - s_{r+1}). \quad (16)$$

Finally, in line 8 we can compute V' in time $O(1)$ given its previous value using the identity (true unless j is very close to r or k is very close to K)

$$\begin{aligned} & V_3(\{j+1, \dots, r\} \cup \{k, \dots, K\}) \\ &= V_3(\{j+1, \dots, r\} \cup \{k+1, \dots, K\}) \\ &\quad + e_k V_2(\{j+1, \dots, r\} \cup \{k+1, \dots, K\}) \\ &= V_3(\{j+1, \dots, r\} \cup \{k+1, \dots, K\}) \\ &\quad + e_k (V_2(\{j+1, \dots, r\}) + t_{k+1} + (s_{j+1} - s_{r+1})s_{k+1}) \end{aligned}$$

and the identity (16) (with $j+1$ in place of $j+2$). The simple but numerous special cases (signalled by “unless”) should be considered separately.

8 When can we expect the U_2 merging function to be effective?

This section is an attempt of a theoretical explanation of the efficiency of the U_2 merging function in our simulation and empirical studies. We start from an alternative representation of U_2 , which will shed some light on its expected performance.

Let $\mathbf{e} := (e_1, \dots, e_K) \in [0, \infty)^K$, $M_1 = U_1(\mathbf{e})$ be the arithmetic mean of e_1, \dots, e_K ,

$$M_2 := \sqrt{\frac{e_1^2 + \dots + e_K^2}{K}}$$

be the quadratic mean of e_1, \dots, e_K , and

$$\text{var}(\mathbf{e}) := \frac{1}{K} \sum_{k=1}^K (e_k - M_1)^2 = M_2^2 - M_1^2$$

be the sample variance of e_1, \dots, e_K .

Lemma 8.1. *For any \mathbf{e} ,*

$$U_2(\mathbf{e}) = M_1^2 - \frac{1}{K-1} \text{var}(\mathbf{e}). \quad (17)$$

Proof. By definition,

$$\begin{aligned} U_2(\mathbf{e}) &= \frac{1}{K(K-1)/2} \sum_{i < j} e_i e_j = \frac{1}{K(K-1)} \left(\left(\sum_i e_i \right)^2 - \sum_i e_i^2 \right) \\ &= \frac{K}{K-1} M_1^2 - \frac{1}{K-1} M_2^2 = M_1^2 - \frac{1}{K-1} \text{var}(\mathbf{e}). \quad \square \end{aligned}$$

Corollary 8.2. For any \mathbf{e} ,

$$\text{var}(\mathbf{e}) \leq (K-1)M_1^2.$$

For some $\mathbf{e} \neq 0$ the inequality holds as equality.

Proof. The first statement follows from $U_2(\mathbf{e}) \geq 0$, and an example for the second one is $\mathbf{e} := (K, 0, \dots, 0)$. \square

According to Corollary 8.2,

$$\text{rvar}(\mathbf{e}) := \frac{\text{var}(\mathbf{e})}{(K-1)M_1^2},$$

which we will call the *relative (sample) variance* of \mathbf{e} , is a dimensionless quantity in the interval $[0, 1]$. When $\mathbf{e} = 0$, we set $\text{rvar}(\mathbf{e}) := 0$. The relative variance is zero if and only if all e_i coincide, and it is 1 if and only if all e_i but one are zero.

Using the notion of relative variance, we can rewrite (17) as

$$U_2(\mathbf{e}) = M_1^2(1 - \text{rvar}(\mathbf{e})). \quad (18)$$

We can see that the method of this paper based on U_2 has a potential for improving on the U_1 method of [20], but the best it can achieve is squaring the entries of the discovery matrix. An entry of a discovery matrix based on U_1 is squared when we replace U_1 by U_2 if the multiset of e-values on which the min in (7) is attained (for $F := U_1$) consists of a single value. Otherwise the discovery matrix based on U_2 suffers as the e-values become more diverse.

The identity (18) shows that $U_2 > M_1$ (i.e., U_2 improves on $U_1 = M_1$) if and only if

$$\text{rvar}(\mathbf{e}) < 1 - \frac{1}{M_1}.$$

For example, if the result of applying U_1 is borderline strong evidence against the null hypothesis, $U_1 = 10$, it is improved by U_2 if and only if $\text{rvar}(\mathbf{e}) < 0.9$. For comparison:

- in our simulation studies, the relative variance of the whole set of 200 likelihood ratios (10) (those used in Figure 1) is approximately 0.24, and the relative variance of the 20 largest of them is approximately 0.23;
- in our empirical studies, the relative variance of the whole set of 6033 Monte Carlo e-values (12) (those used in Figures 3 and 4) is approximately 0.035; the relative variance of the 200 largest among those 6033 e-values is approximately 0.031, while the relative variance of the 20 largest is approximately 0.028.

9 Conclusion

This paper has given examples of se-merging functions, namely U_n , which can be successfully applied, in simulation and empirical studies, to multiple hypothesis testing with independent base e-values. An interesting question is whether admissible ie-merging functions different from U_n and their convex combinations can also be useful for this purpose. Such functions definitely exist; e.g., in [19, Remark 4.3] we show that

$$f(e_1, e_2) := \frac{1}{2} \left(\frac{e_1}{1+e_1} + \frac{e_2}{1+e_2} \right) (1 + e_1 e_2) \quad (19)$$

is an admissible ie-merging function. To check that $f(e_1, e_2)$ does not have the form $ae_1 + be_2 + ce_1e_2$, it suffices to set $e_2 := 0$. We can even show that (19) is not an se-merging function: see [21, Example 2].

In our computational experiments in Sections 5 and 6 we only study the case of independent base e-values, while in Sect. 4 we study the more general case of sequential e-values. The reason is that, while our main procedure works perfectly well in the sequential case from the mathematical point of view, it appears to be natural from the practical point of view only in the independent case. In the sequential case, as described in Sect. 4, it is natural to allow testing the same hypothesis more than once. Another interesting direction of further research is to design such more general sequential procedures.

In Sect. 7 we described an $O(K^2)$ algorithm implementing our procedure for computing one row of discovery matrices based on U_n . It is an open problem to find $O(K)$ procedures for computing one row of discovery matrices based on U_n with $n > 1$ (for $n = 1$ it was done in [19, Proposition 4.1]), or to prove that such procedures do not exist.

Acknowledgments

We are grateful to Yuri Gurevich for useful discussions. We also thank the editors and two anonymous referees for helpful comments about the journal version of this paper. In our simulation and empirical studies we used Python and R, including the package `hommel` [9].

V. Vovk's research has been partially supported by Astra Zeneca, Stena Line, and Mitie. R. Wang is supported by the Natural Sciences and Engineering Research Council of Canada (RGPIN-2018-03823, RGPAS-2018-522590).

References

- [1] Daniel J. Benjamin and James O. Berger. Three recommendations for improving the use of p-values. *American Statistician*, 73(Supplement 1):186–191, 2019.
- [2] Daniel J. Benjamin, James O. Berger, Magnus Johannesson, Brian A. Nosek, Eric-Jan Wagenmakers, Richard Berk, Kenneth A. Bollen, et al.

Redefine statistical significance: We propose to change the default p-value threshold for statistical significance from 0.05 to 0.005 for claims of new discoveries (Comment). *Nature Human Behaviour*, 2:6–10, 2018.

- [3] Ziyu Chi, Aaditya Ramdas, and Ruodu Wang. Multiple testing under negative dependence. Technical Report [arXiv:2212.09706 \[math.ST\]](https://arxiv.org/abs/2212.09706), [arXiv.org](https://arxiv.org/) e-Print archive, May 2023.
- [4] Bradley Efron. *Large-Scale Inference: Empirical Bayes Methods for Estimation, Testing, and Prediction*, volume 1 of *IMS Monographs*. Cambridge University Press, Cambridge, 2010.
- [5] Bradley Efron and Trevor Hastie. *Computer Age Statistical Inference: Algorithms, Evidence, and Data Science*, volume 5 of *IMS Monographs*. Cambridge University Press, Cambridge, 2016.
- [6] Ronald A. Fisher. *Statistical Methods for Research Workers*. Oliver and Boyd, Edinburgh, fourth edition, 1932. Section 21.1 on combining independent p-values first appears in this edition and is present in all subsequent editions.
- [7] Ronald A. Fisher. *Statistical Methods and Scientific Inference*. Hafner, New York, third edition, 1973.
- [8] Christopher R. Genovese and Larry Wasserman. A stochastic process approach to false discovery control. *Annals of Statistics*, 32:1035–1061, 2004.
- [9] Jelle J. Goeman, Rosa Meijer, and Thijmen Krebs. *hommel: Methods for closed testing with Simes inequality, in particular Hommel’s method*, 2019. R package version 1.5, available on CRAN.
- [10] Jelle J. Goeman, Rosa J. Meijer, Thijmen J. P. Krebs, and Aldo Solari. Simultaneous control of all false discovery proportions in large-scale multiple hypothesis testing. *Biometrika*, 106:841–856, 2019.
- [11] Jelle J. Goeman and Aldo Solari. Multiple testing for exploratory research. *Statistical Science*, 26:584–597, 2011. Correction: 28:464.
- [12] Peter Grünwald, Rianne de Heide, and Wouter M. Koolen. Safe testing. Technical Report [arXiv:1906.07801 \[math.ST\]](https://arxiv.org/abs/1906.07801), [arXiv.org](https://arxiv.org/) e-Print archive, March 2023. To appear in the *Journal of the Royal Statistical Society B* (with discussion).
- [13] Harold Jeffreys. *Theory of Probability*. Oxford University Press, Oxford, third edition, 1961.
- [14] Samuel Kotz, N. Balakrishnan, Campbell B. Read, Brani Vidakovic, and Norman L. Johnson. Null hypothesis. In *Encyclopedia of Statistical Sciences*, volume 8, pages 5683–5684. Wiley, Hoboken, NJ, second edition, 2006.

- [15] Glenn Shafer. The language of betting as a strategy for statistical and scientific communication. *Journal of the Royal Statistical Society A*, 184:407–478, 2021.
- [16] R. John Simes. An improved Bonferroni procedure for multiple tests of significance. *Biometrika*, 73:751–754, 1986.
- [17] Dinesh Singh, Phillip G. Febbo, Kenneth Ross, Donald G. Jackson, Judith Manola, Christine Ladd, Pablo Tamayo, Andrew A. Renshaw, Anthony V. D’Amico, Jerome P. Richie, Eric S. Lander, Massimo Loda, Philip W. Kantoff, Todd R. Golub, and William R. Sellers. Gene expression correlates of clinical prostate cancer behavior. *Cancer Cell*, 1:203–209, 2002.
- [18] Vladimir Vovk, Alex Gammerman, and Glenn Shafer. *Algorithmic Learning in a Random World*. Springer, Cham, second edition, 2022.
- [19] Vladimir Vovk and Ruodu Wang. E-values: Calibration, combination, and applications. *Annals of Statistics*, 49:1736–1754, 2021.
- [20] Vladimir Vovk and Ruodu Wang. Confidence and discoveries with e-values. *Statistical Science*, 38:329–354, 2023. Revised arXiv version: [arXiv:1912.13292 \[math.ST\]](https://arxiv.org/abs/1912.13292) (November 2022).
- [21] Vladimir Vovk and Ruodu Wang. Merging sequential e-values via martingales. *Electronic Journal of Statistics*, 18:1185–1205, 2024.
- [22] Ruodu Wang and Aaditya Ramdas. False discovery rate control with e-values. *Journal of the Royal Statistical Society B*, 84:822–852, 2022.
- [23] Ronald L. Wasserstein, Allen L. Schirm, and Nicole A. Lazar. Moving to a world beyond “ $p < 0.05$ ”. *American Statistician*, 73(Supplement 1):1–19, 2019.