**SILR: A New Exact Test for Demonstrating
That an Effect Exists in Binary Trials**

Richard B. Darlington

Department of Psychology, Cornell University

204 Fairmount Avenue, Ithaca NY 14850

rbd1@cornell.edu

Zhengze Zhou

Department of Statistics and Data Science, Cornell University

301 Malott Hall, Ithaca, New York, 14850

+1(607)319-6302

zz433@cornell.edu

Abstract

Suppose N experimental participants each undergo t binary trials, and you want to show that at least some of the participants have higher success rates than would be expected if the true success rate for all participants on all trials were the same known value p_{null} . Under broad circumstances, a new test called SILR has substantially higher power than six other tests that might be used for this purpose. In one example, SILR yielded a significance level of 0.0032 while the other six tests all yielded values above 0.05. In one power analysis, SILR's power exceeded that of the standard binomial test, applied to the results of all trials for all participants, even when the sample size for SILR was less than 30% that for the binomial.

We provide a free R program for SILR. It can also find confidence limits on the number of experimental participants whose true hit rate exceeds the null value, and on the highest and lowest true hit rates of any participants. SILR does not show conclusively that any particular participant's true hit rate differs significantly from the null rate. If such tests are desired, that can be done with Bonferroni-corrected binomial tests on individual participants. But SILR often has much higher power than these tests.

Keywords

beta binomial

binomial

binary tests for ability or preference

Introduction

What SILR is for

Can any dogs smell cancer? Can anyone read the thoughts of someone in the next room? Can anyone reliably identify lies by mere observation, from the speaker's facial musculature or tone of voice? Can anyone reliably distinguish sugar from artificial sweeteners? Does anyone *prefer* artificial sweeteners? We'll call these questions *existence* questions, since they ask whether a certain ability or preference exists at all. Food science is one area of psychology in which questions like this arise especially often. We actually created SILR in response to a statistical question from a psychologist/food scientist, and three of our later references are to works in food science.

These questions can be studied with experiments in which each of N participants undergoes t trials, where $t > 1$, in which the probability of each success or hit is a known probability p_{null} . This paper describes a new statistical test for these questions, and shows that this test is often more powerful for that purpose than six competing tests. The new test is called SILR because it's based on single-individual likelihood ratios – that is, likelihood ratios calculated for single individuals. The test is quite complex internally, but a free R program typically runs in a few seconds or less. Access to the program is described at the end of section *SILR*.

We assume that each experimental participant i has an unknown true long-term hit rate p_{true_i} , and that these values may differ across participants. The highest of these N unknown values is defined as p_{max} and the lowest as p_{min} . Also let n_{pos} denote the unknown number of experimental participants for whom $p_{true_i} > p_{null}$, and let n_{neg} denote the number for whom

$p_{true_i} < p_{null}$. Our computer program can also find, for any confidence level the user specifies, lower confidence limits on n_{pos} , n_{neg} , and p_{max} , and an upper confidence limit on p_{min} .

We describe an *upper-tail* test which could be used to show $p_{max} > p_{null}$ or equivalently that $n_{pos} > 0$. A *lower-tail* test can show that $p_{min} < p_{null}$ or equivalently that $n_{neg} > 0$. All the questions in our opening paragraph would be studied with upper-tail tests. A significant result in a lower-tail test might mean that some participants have a consistently mistaken approach to the problem, or are deliberately trying to fail. Or if the experiment is testing preferences between stimuli A and B, it could mean merely that some participants prefer B over A rather than A over B. We also describe a *some-tail* test, which can show that at least one of the first two effects exists. That test might be the most powerful of the three if $p_{true_i} > p_{null}$ for some participants and $p_{true_i} < p_{null}$ for others, but would otherwise be less powerful than the others.

Researchers may or may not be immediately interested in showing that particular experimental participants deviate from the null. In testing whether any dogs can smell cancer, they may have no specific plans to use any of the experimental dogs as future cancer detectors. Or if the question is whether any human participants can detect a difference in taste when a food is made with either cheaper or more expensive ingredients, there may be no interest in knowing which individual participants can do this. Conclusions about individual participants can be reached via Bonferroni-corrected binomial tests on individuals, though not by SILR alone. But the power of SILR may far exceed that of the Bonferroni tests, so SILR may be useful even when researchers would ultimately like to study individual participants.

For an upper-tail test, SILR uses each participant's number of hits to compute a binomial likelihood ratio (LR) testing the null hypothesis that $p_{true_i} \leq p_{null}$ for that one participant. The

product of these N values of LR is the LR for the hypothesis $p_{max} = p_{null}$. Similar points apply for lower-tail tests. The natural log of the whole-sample LR is SILR's test statistic G . Let g_i denote the natural log of the LR computed for participant i . Then $G = \sum g_i$. Thus, computing G is quite simple. These procedures are described in detail early in section *SILR*.

In the computer program, the user specifies an upper-tail, lower-tail, or some-tail test, and enters p_{null} , a significance level α , and the number of participants with each possible number of hits. The program then returns all the values described above. The program does not perform upper-tail, lower-tail, and some-tail tests simultaneously, because doing so might tempt a user to emphasize *post hoc* whichever test came out most significant. To minimize the multiple-test problem, the user should decide which of these tests is of primary interest.

SILR is a randomization test as defined by Edgington (2011). These tests use some sort of randomization – typically random ordering of stimuli, or random division of experimental participants into two or more groups – and this randomization allows us to compute a probability (under the null hypothesis) for every possible outcome. Then an experiment's significance level is the sum of the probabilities of all the possible outcomes in which the effect size is at least as large as the observed outcome. These tests can show that observed effects could not reasonably be caused by chance, even without assuming that participants were drawn randomly from a larger population; SILR shares that advantage. Edgington and Onghena (2007, 6-17) describe the long and important history of randomization tests in modern statistics. Sir Ronald Fisher didn't use that term, but repeatedly advocated those tests; see Fisher (1925, 1935, 1936). In fact, Fisher (1936, p. 59) justified ordinary parametric tests primarily as approximations to randomization tests, because genuine random sampling from larger populations is so rare in the real world.

Six alternatives to SILR

This section describes six statistical methods which might seem to provide reasonable alternatives to SILR, but the power of all six is often far below that of SILR. This gain in power occurs because of SILR's sensitivity to what we call *hit concentration* – a tendency of hits to be concentrated in a small fraction of the participants. For instance, suppose $p_{null} = 0.5$, $N = 37$ and $t = 5$. Thus, there are $5 \cdot 37$ or 185 trials altogether, of which we would expect half, or 92.5, to be hits. If 104 hits are observed, a one-tailed binomial test yields a significance level ps of 0.053, just missing the 5% level. But if those 104 hits are concentrated heavily in just a few participants, intuition suggests the data are less consistent with the null than if the same number of hits were spread out more evenly across all participants. The binomial test ignores that difference and thus ignores hit concentration. But SILR considers hit concentration, and may yield a highly significant result in this example, as shown when we return to it at the end of this section.

Let TBB, for “trial-based binomial,” denote the standard binomial test just mentioned. TBB is actually the most powerful possible test if all participants have the same true long-term hit rate p_{true_i} . But TBB loses power relative to SILR as hit concentration increases, since the test statistic in TBB is effectively just the mean hit rate. Another test using the mean hit rate is the t test on the mean number of hits per participant. But just the opposite of SILR's sensitivity to hit concentration, the t test actually becomes *less* significant if the participants' numbers of hits become more diverse as their mean stays constant. The beta-binomial test, discussed by Bi (2006, pp. 106-137) and Lawless and Heymann (2010, pp. 96-97), might appear to be a more serious competitor to SILR, since both those tests are designed specifically for experiments in which each participant undergoes two or more binary trials. But like the t test, it uses the mean hit rate as

its test statistic and its result actually becomes less significant as the participants become more diverse in their numbers of hits while their mean stays constant.

A fourth alternative was mentioned in the previous section: use the classic binomial test to test each participant separately, select the one participant with the smallest p_s , and apply a Bonferroni correction to correct for the post hoc selection of that one participant. This approach is practical if t is high enough, and p_{null} is low enough, to allow highly significant results for single participants. E.g., if $t = 5$ and $p_{null} = 0.2$, then before a Bonferroni correction a single-participant significance level may be as low as 0.2^5 or 0.00032. Since $0.05/0.00032 = 156.25$, after correction that result would still be significant at the .05 level even if that one participant were the best performer in a sample up to size 156. And unlike SILR, this test allows investigators to say *which* participants are clearly performing better than chance. But this approach is useless for many other combinations of p_{null} and t . For instance, if $p_{null} = 0.5$ and $t = 5$, the lowest possible one-tailed single-individual probability is 0.5^5 or 0.03125, which wouldn't survive a Bonferroni correction even if alpha were 0.05 and N were only 2. And even when the approach isn't completely useless, its power is likely to be far below that of SILR, since the Bonferroni approach uses so little of the data, ignoring for instance the fact that there may be several participants with many hits.

A fifth alternative is what we call a double-binomial (DB) test, used by Wajid and Halpern (2012). This test counts the number of participants who reach or exceed a certain required number of hits r , uses the binomial distribution to calculate the probability pr that any one participant will do so, and then uses the binomial distribution a second time, this time with pr as the null probability, to find the probability that the number of participants with r or more hits

will equal or exceed the number observed to meet that criterion. For instance, if we had $p_{null} = 0.2$, $t = 5$, and $r = 3$, we would have $pr = \text{prob}(3 \text{ hits}) + \text{prob}(4 \text{ hits}) + \text{prob}(5 \text{ hits}) = [5!/(3!2!)] \cdot 0.2^3 \cdot 0.8^2 + [5!/(4!1!)] \cdot 0.2^4 \cdot 0.8 + 0.2^5 = 0.05972$. For any given experiment, DB is actually a family of tests, since r can range anywhere from 1 to t . In naming the test we'll put r in a subscript, such as DB_3 or DB_t , where DB_t is the test counting the number of participants who succeed on all trials. There is a rare circumstance where DB_t would clearly be the best test. That's the case in which the population has only two subsets: those who respond according to chance, and those who respond perfectly. In most experiments the chance of a perfect response pattern (t hits) is very small under the null, so a binomial test on the number of such participants would be very powerful if that number is at all large.

A final alternative to SILR is a chi-square test which compares the observed frequency of each possible number of hits to the frequency expected under the null. The higher t , the more power this test loses by having to allow for patterns which could occur by chance but which are not expected under either the null or alternative hypotheses, such as the pattern (when $t = 5$) in which many participants have 1, 3, or 5 hits but few have 0, 2, or 4 hits. That's one of many patterns which, under any reasonable hypothesis, are not expected except by chance. No reasonable test would treat such patterns as if they were fully consistent with the null. But the most powerful tests would give such patterns intermediate values of the test statistic, reserving the highest values of the test statistic for those patterns which are both inconsistent with the null and highly consistent with the hypothesized alternative. That's the essence of the generally-accepted likelihood-ratio principle, explained more fully in the next section. In problems of the sort we're considering, the chi-square test essentially ignores this principle, especially when t is

high, and loses power as a result. The chi-square test is omitted from the later section on *Power Analysis*, both for this reason and because it's only an approximate test, and there's little point in comparing the power of exact tests to that of an approximate test.

We return now to the example at the beginning of this section, in which $N = 37$, $t = 5$, $p_{null} = 0.5$. Suppose the number of participants with each number of hits from 0 to 5 is 1, 5, 10, 10, 6, 5; that produces the 104 hits in that example. If we had just 32 participants and their numbers of hits exactly fitted the null hypothesis, our list of frequencies for 0 to 5 hits would be 1, 5, 10, 10, 5, 1. Thus our 37-case outcome is highly consistent with the possibility that 32 of our 37 participants fit the null while 4 others get 5 hits each and one gets 4 hits; that would bring the six frequencies to exactly those observed for the 37 cases. Thus, these data suggest a substantial amount of hit concentration. For this outcome, the TBB, t , beta-binomial, Bonferroni, DB_4 , and chi-square test all yield $ps > 0.05$ while SILR yields $ps = 0.0032$, which is below $1/15$ of 0.05. Another test, DB_5 , also yields a highly significant result ($ps = 0.0056$), since that's the probability of finding 5 or more participants in this example with 5 hits each. But before seeing the data we would have been unlikely to choose test DB_5 over DB_4 , and of course we can't choose a test *because* it gave a significant result while others didn't.

This was just one example, but the later section on *Power analysis* gives a more comprehensive power comparison of tests SILR, TBB, and DB. The latter two tests were selected for that analysis because we saw earlier that TBB is the most powerful possible test if all values of p_{true_i} are equal, falling somewhere between p_{null} and 1.0. And DB_t is the most powerful possible test in just the opposite case, with all values of p_{true_i} falling either at p_{null} or at 1.0. Our

power analysis compares the power of these tests to SILR in 200 other cases which seem more likely to arise in real data.

Notation

The following notations are introduced one at a time throughout the article. They're collected here as a memory aid. A few notations, which are used only in the same paragraph in which they're introduced, are omitted from this list.

α : a chosen level of significance, such as 0.05

DB_r : "double binomial" test family which uses the binomial distribution twice

$g_i = \ln(lr_i)$: a measure of inconsistency between the null H and the data of participant i

$G = \sum g_i$ = test statistic for SILR

h_i : number of hits or other focal responses for participant i

lr_i = likelihood ratio for participant $i = pc_i/p0_i$

LR = whole-sample likelihood ratio = $\prod lr_i$, where $\prod$ denotes product

N : number of participants

$nullprop$: the proportion of the population for which $p_{true_i} = p_{null}$. Used in power analyses.

p_j = probability of outcome j in a multinomial experiment

$p0_i$: probability under the null that participant i achieves exactly h_i hits

pc_i : probability that participant i achieves exactly h_i hits under the H most consistent with that number

p_{null} : null probability of the focal response on a single trial

p_{true_i} : true single-trial probability of the focal response for participant i

r : number of hits constituting a “successful” response for one participant in test family DB_r

ri : rounding interval. Values of g are rounded to the nearest multiple of ri

SILR: the new test introduced here, “single-individual likelihood ratio”

t : number of trials per participant

TBB: the classic binomial test which treats all $N \cdot t$ trials as independent

SILR

The Neyman-Pearson lemma (Mood et al, 1974, pp. 411-412) says that under broad conditions, the most powerful tests are those based on likelihood ratios. Since we are allowing the values of p_{true_i} to vary across participants, the most relevant likelihood ratios are those calculated for each participant separately; each participant’s ratio is lr_i . Our discussion is simplest if we focus on upper-tail tests; the comparable conclusions for lower-tail and some-tail tests are generally obvious. We can find lr_i for any participant i by first using the binomial distribution to calculate the probability $p0_i$ of exactly h_i hits under the null. Then use the binomial again to calculate the probability pc_i of exactly h_i hits if $p_{true_i} = h_i/t$, since this is the value of p_{true_i} which maximizes the chance of exactly h_i hits. We define $lr_i = pc_i/p0_i$. (Here we deviate from the usual practice of putting $p0_i$ on top. It doesn’t affect the final outcome, and this approach eliminates negative logarithms in later calculations.) If participant i ’s number of hits is equal to or less than the number expected under the null, we needn’t bother to perform these calculations because we will always find $lr_i = 1$. As described in our opening subsection, we denote the N individual values of $\ln(lr_i)$ as g_i , and denote $\ln(LR)$ as G , with $G = \sum g_i$. G is our final test statistic.

For the SILR test we need to find the distribution of G under the null hypothesis. This paragraph briefly describes three approaches we considered for this purpose, but abandoned in favor of the convolution approach described in later paragraphs. In the first of these, we used the multinomial theorem to compute the probability of each of the $(N+t)!/(N!t!)$ possible outcomes of an SILR study, and ranked these probabilities by their G values to find the distribution of G . That approach is quite practical for smaller values of N and t . But for larger problems, even a fast computer would take too long. For instance, if $N = 100$ and $t = 10$ then $(N+t)!/(N!t!) = 4.7\text{e}+13$. Thus, if we could process one million outcomes per second, the analysis would take almost 18 months. The second and third approaches were to use the normal or beta distribution to approximate the distribution of G . The normal distribution seems especially inappropriate since G 's distribution is so skewed, and even the beta distribution seems less accurate than the convolution approach, which seems satisfactory in all major respects.

We now describe the convolution approach, which seems superior to the three alternatives just mentioned. Each user of SILR sets the value of a "rounding interval" ri , and rounds all values of g_i to the nearest multiple of ri . Setting ri larger makes the later calculations go faster, while setting it lower presumably raises the power of the SILR test by making it more closely approximate a true likelihood-ratio test. We have found that setting ri to 0.01 achieves very satisfactory values of both speed and power. So, in our computer program, ri is 0.01 by default, though users who wish may reset it. G is redefined as the sum of these rounded values of g_i .

The SILR computer program first uses the binomial distribution to find the probability of every possible number of hits for a single participant under the null hypothesis, and computes

the rounded value of g_i for each of those numbers of hits. The program then constructs a *single participant convolution vector* SPCV. A vector is a row or column of numbers. Each cell in the SPCV vector has a *label* which is not the same as the value in the cell. The labels run in increments of r_i , from 0 up to $maxg$, the highest value of g_i that was found. Thus, the number of cells in SPCV is $maxg/r_i + 1$. The extra 1 is because the first cell is labeled 0. So SPCV has a cell for every possible value of g_i . We enter each of the aforementioned binomial probabilities into the cell of SPCV whose label equals the g_i corresponding to that probability. Sometimes two or more of those probabilities will have the same value of g_i ; those probabilities are then summed within the appropriate cell. That completes the construction of SPCV. The values in SPCV always sum to 1, and typically most of those values are 0.

The SILR program uses the fact that any two vectors can be *convolved* to form a new *convolution vector*. This process is described in Casella and Berger (2002) and other sources, and repeated here. Suppose we have two vectors rx and ry with nx and ny cells respectively. We can first form a nx by ny *product matrix* in which each cell ij equals $rx_i \cdot ry_j$. Table 1 shows an example. The bold numbers in the first column show rx , and those in the bottom row show ry . The 3 x 4 product matrix contains the unbold numbers in the center of the table, surrounded by a bold frame.

If i indexes the cells in rx and j indexes the cells in ry , then $i+j$ can range from 2 to $nx+ny$, so the number of different possible values of $i+j$ is $nx+ny-1$. To convolve rx with ry we set up a new empty convolution vector with that number of cells. Then, for every cell ij in the product matrix, we add the value in that cell to cell $i+j-1$ of the convolution vector. All the cells with the same value of $i+j-1$ fall in the same upward-sloping diagonal of the product matrix, so each cell

in the convolution vector equals the sum of the entries in one of those diagonals. Those sums appear in ***bold italics*** in the top row and final column of Table 1, so that's the convolution vector, which in Table 1 has its own bold frame. In SILR the entries in any convolution vector always sum to 1, and that's true here.

Table 1. Computation of a simple convolution vector, shown in ***bold italics***

<i>rx</i>		<i>0.03</i>	<i>0.14</i>	<i>0.29</i>	<i>0.32</i>
<i>0.3</i>	0.03	0.09	0.12	0.06	<i>0.18</i>
<i>0.5</i>	0.05	0.15	0.20	0.10	<i>0.04</i>
<i>0.2</i>	0.02	0.06	0.08	0.04	
<i>ry</i>	<i>0.1</i>	<i>0.3</i>	<i>0.4</i>	<i>0.2</i>	

Convolution commands are available in many mathematical and statistical packages, for instance as *convolve* in R. In SILR the vectors to be convolved may each have thousands of cells, so a product matrix would have millions of cells and consume large amounts of computer memory. But there is no need to construct that matrix; instead set up an empty convolution vector, then use nested loops to compute the products one at a time, and add each one immediately to cell $i+j-1$ of that vector. The product matrix was introduced here only to make it easier for readers to see and check every step.

Each cell of SPCV contains the probability (under the null hypothesis) that any single participant has a value of g_i equal to that cell's label, so that vector gives a complete description of the null distribution of the values of g_i for any one experimental participant. Suppose we wanted a similar vector describing the null distribution of G for a sample of 2 participants, whom

we assume to be mutually independent. Suppose we convolve SPCV with itself. Then each of the diagonals in the product matrix will contain all the probabilities associated with a particular value of G , so the sum within each diagonal is the probability of that value of G for a sample of size 2. Thus, the vector we want is simply the convolution of SPCV with itself. Call that vector CV2, and let CV1 be another name for SPCV. By the same argument the distribution of G for a sample of size 3 is CV3, the convolution of CV1 with CV2. The distribution of G for a sample of size 4 is CV4, which can be found either by convolving CV3 with CV1 or by convolving CV2 with itself. By similar logic we can construct a convolution vector showing the null distribution of G for any sample size N we wish. The fastest way to do that would be to execute all the doubling convolutions possible without exceeding N , while retaining in RAM all those vectors as we go. Then, after passing $N/2$, at each step we add the largest retained vector which doesn't put us past N . For instance, if $N = 27$ we find CV2 then CV4 then CV8 then CV16. Then we add CV8 to get CV24, then CV2 to get CV26, then CV1 to get CV27.

Table 2 illustrates this process. The amount of calculation in the process increases rapidly with N and t , so we give an example with very small values. Consider an upper-tail problem with $N = 5$, $t = 2$, $p_{null} = 0.2$, $r_i = 0.01$. Suppose 3 participants get no hits and 2 get 2 hits each. That's consistent with the possibility that 2 of the participants have high values of p_{true} while the other 3 are guessing. If 3 participants guess randomly, the probability they will all get 0 hits is 0.8^6 or 0.262, so these imaginary results are quite consistent with that premise. When $p_{null} = 0.2$ and $t = 2$, the rounded g -values for participants with 0, 1, and 2 hits are respectively 0, 0.45, and 3.22, so we have $G = 3 \cdot 0 + 2 \cdot 3.22 = 6.44$ for this 5-case example. Table 2 shows the nonzero entries in the vectors p_1 , p_2 , p_4 , and p_5 we would get from the convolutions for 5 participants. Because

the G associated with each cell in a convolution vector starts with 0 and in this example increases by 0.01 for each cell, the cell associated with the observed G for all 5 cases is cell $100G+1$ in $p5$. Since $G = 6.44$, that's cell 645. Thus, ps for this problem is the sum of all entries in $p5$ starting with that cell. That sum rounds to 0.0147, so $ps = 0.0147$. For this problem the standard binomial test notes there were 4 hits in 10 trials, so the binomial ps is the probability of 4 or more hits. With $pnull = 0.2$, that value is 0.1209, over 8 times as large as the SILR value of ps .

Vector $p2$ was constructed by first forming a 323×323 product matrix $p2m$ in which $p2m[i,j] = p1[i] \cdot p1[j]$. We can express $p2m$ compactly if we delete every row and every column which is all zeros. That gives us the 3×3 matrix in Table 3. In that table, the bold italicized numbers are the row numbers i and column numbers j remaining in $p2m$, the other bold integers are the values of $i+j-1$ for each cell, and the non-integers are the nonzero values in $p2m$. We form the convolution vector CV2, with $323+323-1$ or 645 cells, by adding each non-integer in Table 3 to cell $CV2[i+j-1]$. After those additions, $CV2[46] = 0.2048 + 0.2048 = 0.4096$, which appears in column $p2$ of Table 2. The other probabilities in column $p2$ are calculated similarly from entries in Table 3. Values in columns $p4$ and $p5$ of Table 2 could be calculated similarly, from other tables like Table 3.

Table 2. Nonzero entries in vectors $p1$, $p2$, $p4$, and $p5$, for an upper-tail SILR problem with $p_{null} = 0.2$, $t = 2$, when 3 participants have no hits, none have 1 hit, and 2 have 2 hits each. Integers within each column are cell numbers within the vectors; non-integers are the contents of those cells. Non-integers are probabilities; they sum to 1 in each column. All these probabilities are exact values.

p1	p2	p4	p5
1 0.64	1 0.4096	1 0.16777216	1 0.1073741824
46 0.32	46 0.4096	46 0.33554432	46 0.2684354560
323 0.04	91 0.1024	91 0.25165824	91 0.2684354560
	323 0.0512	136 0.08388608	136 0.1342177280
	368 0.0256	181 0.01048576	181 0.0335544320
	645 0.0016	323 0.04194304	226 0.0033554432
		368 0.06291456	323 0.0335544320
		413 0.03145728	368 0.0671088640
		458 0.00524288	413 0.0503316480
		645 0.00393216	458 0.0167772160
		690 0.00393216	503 0.0020971520
		735 0.00098304	645 0.0041943040
		967 0.00016384	690 0.0062914560
		1012 0.00008192	735 0.0031457280
		1289 0.00000256	780 0.0005242880
			967 0.0002621440
			1012 0.0002621440
			1057 0.0000655360
			1289 0.0000081920
			1334 0.0000040960
			1611 0.0000001024

Table 3. Remains of product matrix $p2m$ after all empty rows and columns have been deleted.

With directions in the text, this matrix could be used to compute column $p2$ of Table 2.

$i \backslash j$	1	46	323
1	0.4096, 1	0.2048, 46	0.0256, 323
46	0.2048, 46	0.1024, 91	0.0128, 368
323	0.2056, 323	0.0128, 368	0.0016, 645

Convolutions can be sped up by an operation we call *vector trimming*. If we have found a particular value of G in our sample, we don't need to find the separate individual probabilities of all the possible values of G which exceed that observed value. We can instead merge all those probabilities, and the probability of finding exactly the value of G we observed, into a single event of finding a value of G equal to or higher than the value we observed. In the example of Table 2 we would merge the last 6 cells of column $p4$ into a single cell with a value of 0.00909568, and merge the last 10 cells of column $p5$ into a single cell with a value of 0.0147579904. In larger problems this process can reduce computing time substantially.

Convolutions can also be sped up enormously by the Fast Fourier Transform FFT. It was first popularized by Cooley and Tukey (1965), though historical research (Goldstine, 1977) has shown that a version of it was invented by Gauss around 1805. For large problems (the ones most in need of speedup), it can speed up convolutions over 1000-fold, reducing computing times from hours to seconds. Space prohibits an explanation of FFT here. Beyene (2001) discusses in detail the use of FFT for statistical inference, though with no mention of what we call "doubling" or

“vector trimming.” Web searches using “FFT” and “statistical inference” yield many other relevant works. Many statistical packages including R contain the necessary commands for using FFT. Our SILR program uses FFT. We have found that FFT can theoretically introduce some rounding error, but too little to be of any concern in real problems, since it appears mainly when $\alpha < 1e-12$.

We have implemented SILR in R. All relevant codes are posted at <https://github.com/ZhengzeZhou/SILR>. Readers familiar with R and Shiny may directly download the R scripts from the link provided and run the program locally. Others may run `"runGitHub("SILR", "ZhengzeZhou")"` in the console of R studio. Here "SILR" is the repository name, and "ZhengzeZhou" is the username for Github. Readers may also directly access the Shiny program at <https://zhengzez.shinyapps.io/SILR-master/>. Readers who want to try this program, but who have not yet read our later sections on confidence limits for *npos*, *nneg*, *pmax*, and *pmin*, should tell the program not to calculate confidence limits, since they can sometimes slow the program down.

Power Analysis

Design of the power analyses

When comparing the power of SILR to that of TBB and DB, its primary competitors, we used a power-comparison method which is exact in two senses and approximate in a third. It computes the exact power of each test rather than relying on simulations, which are always approximate because of their finite sample sizes. And it uses exactly the same alternative hypothesis H_1 for each test being compared. But our verbal descriptions of those alternative

hypotheses are sometimes only approximate. We approximate the verbal description by positing a finite number of subsets of participants, with values of p_{true} evenly spaced along some interval. For instance, suppose $p_{null} = 0.2$ and we want to calculate power under the alternative hypothesis that values of p_{true} are distributed uniformly from 0.2 to 1.0. We could approximate that by positing that 1/17 of the participants have each of the 17 p_{true} values 0.20, 0.25, 0.30, ..., 1.0. Or if $p_{null} = 0.2$ and H_1 specifies that half the participants fit H_0 while the other half have p_{true} values distributed evenly from p_{null} up to 1.0, we could approximate that by positing that half the participants have $p_{true} = 0.2$ while the other half are divided evenly among the other 16 values of p_{true} mentioned in the first example. Let q denote the number of subsets we have thus specified, so $q = 17$ in both these examples. We derive the binomial distribution of hits h_i for single individuals within each of those q subsets, and then average those q distributions, using the weights we posited. In the first example we would use an unweighted average, while in the second example we would use a weighted average with one weight of 0.5 and 16 weights of $0.5/16$ or 0.03125. Other distributions such as the normal can be approximated almost as easily as uniform distributions; values of p_{true} could still be evenly spaced, but they would be averaged using unequal weights. Recall that our inferential techniques make no reference to a larger population of participants. Thus, our inferential conclusions apply not to a larger population of individuals, but to a hypothetical population of infinitely many responses from the N participants in the study.

Consider again the conditions required, aside from independence, to use convolutions to compute the distribution of a test statistic: (1) the test statistic must equal the sum of scores computed for each participant, (2) the single-participant scores must have only a finite number

of possible values, evenly spaced along a continuum, with the same spacing for all participants, (3) we must know the probability of each of those single-participant scores; probabilities of 0 for some scores are allowed.

We have just seen how to state an alternative hypothesis H_1 in a way that meets these conditions. We know these conditions hold for convolution-based SILR. They also hold for TBB, since its test statistic is the sum of the h -values (number of hits) for individual participants. They also hold for DB; its test statistic is the number of participants achieving at least a specified number of hits, and we can think of that as the sum of individual scores on a scale scored 1 for participants who meet that condition and 0 for those who don't. Thus, the power analyses in the next section were executed entirely with convolutions, using single-participant distributions constructed using the methods of this section.

Power comparison of SILR, TBB, and DB

In the first case used in the power analyses, values of p_{true_i} are distributed evenly from p_{null} up to 1.0. In the other case, some fraction of the population is at p_{null} and the rest is distributed evenly from p_{null} up to 1.0. We merged these cases together by defining $nullprop$ as the proportion of the population with $p_{true_i} = p_{null}$, and studying $nullprop$ values of 0 to 0.9 in increments of 0.1. We studied p_{null} values of 0.1, 0.2, 0.3, 0.4, and 0.5, and t values of 5, 10, 15, and 20. We studied every combination of these three parameters, thus studying $10 \cdot 5 \cdot 4$ or 200 different cases.

Table 4. Relative power of SILR and TBB in 200 cases, under each of two different distributions of p_{true_i} . Entries show the number of cases in each category.

Values of p_{true_i} range from p_{null} to 1.0; see text for details							
SILR superior				TBB superior			
substantial	noticeable	slight	very slight	very slight	slight	noticeable	substantial
96	1	25	44	29	5	0	0
Values of p_{true_i} are all equal and above p_{null} ; see text for details							
SILR superior				TBB superior			
substantial	noticeable	slight	very slight	very slight	slight	noticeable	substantial
0	0	0	0	10	60	0	130

For each of these 200 cases, we compared SILR to TBB as follows. For each method we found the smallest N needed to achieve power of 0.8 or higher when $\alpha = 0.05$; power of 0.8 was suggested by Cohen (1977) as a “standard” power level comparable to 0.05 for α . For each method, we recorded the exact power achieved once power exceeded 0.8. We’ll say that one method “substantially” outperformed the other if the ratio between the two required sample sizes was 1.2 or greater *and* they differed by 2 or more. This combined rule thus excluded ratios of 1.2 between very small sample sizes such as 6 and 5. We’ll say that one method “noticeably” outperformed the other if their required sample sizes differed by 2 or more but the ratio between them was under 1.2. We’ll say that one method “slightly” outperformed the other

if the required sample sizes differed by just 1. And we'll say that one method "very slightly" outperformed the other if the required sample sizes were equal but the first method had the higher power at that sample size. The top half of Table 4 shows how these 200 cases fell into these mutually exclusive categories. In all the cases of slight or very slight outperformance, both methods required sample sizes of 9 or less. Thus, SILR noticeably or substantially outperformed TBB in every case in which one or both methods required sample sizes of 10 or greater. The highest ratio of necessary sample sizes was 3.39. In that case, $t = 15$, $p_{null} = 0.5$, $p_{true_i} = p_{null}$ for 90% of the participants, the two necessary sample sizes were 51 and 173, and the power of SILR slightly exceeded that of TBB even though SILR's sample size was less than 30% that of TBB's.

Partly to check the programming, we modified the last analysis. We used the same 200 cases, but with the distribution of p_{true} modified to favor TBB over SILR. For each case we now made all N values of p_{true} be one-fourth of the way from p_{null} to 1.0; e.g., if $p_{null} = 0.2$ then all values of p_{true} were 0.4. As expected, TBB now outperformed SILR in all 200 cases, because all values of p_{true} were equal in each case. More detailed results appear in the bottom half of Table 4. The highest ratio of necessary sample sizes was 2.0.

In comparing SILR to the DB family of tests, we start with the fact that a real-world investigator would have t different DB tests to choose among, and would usually have little idea which one is best. The question is therefore whether any DB tests ever outperform SILR so widely that investigators would find it worthwhile to try to overcome this hurdle. To address this question, we defined $\max(\text{DB})$ as the maximum power of the t DB tests for a given problem, using the same sample size already determined to just give SILR power of 0.8 or higher. We compared $\max(\text{DB})$ to the already-computed power of SILR, for the same 200 cases used in the top half of

Table 4. The two power values were never exactly equal. DB won just 31 of these 200 comparisons, just 3 of those 31 power advantages exceeded 0.025, and the largest was 0.0364. SILR's largest power advantage over max(DB) was over 7 times as great at 0.281. Since the power advantages of SILR were both larger and more numerous than those of max(DB), the sum of the former was 29.24 times that of the latter. Across all 200 cases, the mean power advantage was 0.063 in favor of SILR. SILR did so well even though we always compared its power to the highest power of as many as 20 DB tests. We have trouble seeing in these results anything that would encourage further consideration of DB tests for problems like those examined.

Finding Confidence Limits on n_{pos} , n_{neg} , p_{max} , and p_{min}

If researchers have shown that a certain ability exists in some people or animals, they will almost always want to know how many people or animals have that ability, and how great those abilities might be. That means they will want to find lower confidence limits on n_{pos} (the number of experimental participants with the ability) and p_{max} (the maximum long-term hit rate of any of their experimental participants). In other cases, they may want to know the comparable confidence limits for n_{neg} and p_{min} . We focus here on n_{pos} and p_{max} , assuming that the comparable conclusions for n_{neg} and p_{min} will be obvious by analogy. As mentioned in the Introduction, all inferential conclusions from an SILR test pertain to the current set of participants. Any conclusions about any larger population are merely the logically obvious ones – e. g., the maximum ability level in the population cannot be less than the maximum ability level of those in the sample.

Confidence limits on n_{pos}

This section considers how to find a lower confidence limit on n_{pos} , the number of experimental participants in the current sample who have the trait in question. This can be done by sorting the participants in the order of their individual g -values, and dropping the highest- g participants out of the analysis, one at a time, until the result is no longer statistically significant. The SILR user need not drop these cases manually; virtual drops are built into the program. The next two paragraphs show that the lower confidence limit in question is the number who must be dropped to make the result non-significant. For instance, if dropping just one participant turns the result non-significant, then the most conservative hypothesis consistent with the data is that just one participant has the trait in question.

To see why this is true, recall first that SILR does not assume that the experimental participants are drawn randomly from a larger population of possible participants, but does assume that trials are mutually independent both within and between participants. Therefore, our inferential conclusions will pertain only to the present set of participants, not some larger population of participants. Imagine first that k of the N participants possess ability X to an enormous degree, while all others lack it entirely. Then in any sample of trials, the k with the ability will always have the k highest g -values. Thus, they will always be the ones removed when we remove k participants. Thus, when we execute the SILR test after removing those participants, we will always be testing the $N - k$ participants who in fact lack the ability. Thus, over many sets of trials involving those same N participants, the proportion of sets showing significance at level α will exactly equal α , so our test will be a valid test of the null hypothesis that the number with ability X does not exceed k . Call this condition E , for “enormous.”

Now suppose the same k participants have ability X , but to a smaller degree than in condition E . Thus, across many sets of trials involving those same participants, the distribution of observed hit-rates will be lower than before. Thus, the chance of getting a significant result (after removing the top k participants) has to be lower than under condition E , even though there is now a greater chance that the bottom $N - k$ participants will include some who actually have ability X . Thus, the test on these participants will be less likely to be significant than under condition E . Thus, the test is valid. We're saying nothing about the Bayesian posterior probability that a significant sample will contain some participants who actually have ability X , but that's true for any ordinary significance test.

Our SILR program applies these ideas as follows. Of course, if the full-sample ps is non-significant, then the lower confidence limit on $npos$ is 0. Otherwise, the program sorts the participants in the order of their g -values, lowest first. It places these sorted g -values in a column, then cumulates that column one cell at a time to form column $gcum$. Thus, cell $gcum[N-k]$ contains the value of G that would remain if the k highest values of g_i were discarded. An inefficient way to find our confidence limit would be to run a separate SILR analysis for each of these N values of G , to find the smallest of those values which is still significant. Our program uses an iterative process to find that value more efficiently. It computes vector $CV1$ as described earlier, then executes doubling convolutions to find $CV2$, $CV4$, etc. As in section SILR, if another doubling would put it over N , the program executes the largest convolution possible which doesn't put it over N . As in that section, the program holds all these CV vectors in memory, because they may be needed later. But unlike in that earlier section, each new CV is used, together with the matched value of $gcum_u$, to find the corresponding value of ps_u , which tests

the null hypothesis that the u participants with lowest g_i all lack the ability or preference in question. When it finds $ps_u < \alpha$, it gives that value of u the name *uppern*, and gives the largest previous value of u the name *lowern*. It then finds the largest integer u which does not exceed $(uppern+lowern)/2$, and computes the CV and ps_u for that value of u . If it finds $ps_u < \alpha$, then that value of u becomes the new value of *uppern*; otherwise it becomes the new value of *lowern*. Any CV which yields $ps_u > \alpha$ is retained in memory, because it will be used to compute the next CV. This process continues until $uppern = lowern + 1$. The program then stops and reports $N - lowern$ as the lower confidence limit on the number of participants with the ability or preference in question.

Confidence limits on p_{max}

If N is large, this section of the computer program may take much longer than other sections, so the program is written to make this section optional. It uses the value of α entered by the user, and the whole-sample value of ps calculated by the main part of the program. If $ps > \alpha$, there is usually no interest in finding a confidence limit on p_{max} , so this section of the program is not executed at all. Otherwise, the program tries a new value of p_{null} halfway between the previous value and 1.0. It keeps repeating this process until it finds $ps > \alpha$. It then tries a new value of p_{null} halfway between the last two values of p_{null} , whose ps values straddle α . It then repeats this last step as many times as necessary, each time trying a new value of p_{null} halfway between the two previous values of p_{null} which produce the two values of ps which most closely straddle α . This process may take a while if there are a great many participants. Users may sometimes prefer to turn off the confidence limit on p_{max} , and

approximate it manually by successively entering various values of p_{null} until finding one which yields a value of p_s sufficiently near alpha.

Two examples

For an artificial example, suppose each trial has four choices, of which only one is correct, so we set p_{null} at 0.25. Suppose 68 participants have no ability at all to do the task, so they perform at chance. The binomial theorem says the number of those participants getting from 0 to 5 hits will be about (16 27 18 6 1 0). Suppose 34 other participants all have true long-term hit rates of 0.75. The binomial theorem says the comparable list of frequencies for those is (0 1 3 9 13 8), so the list for all 102 participants is (16 28 21 15 14 8). For all 102 participants we find $p_s = 1.15e-24$. If alpha = 0.05 so we're finding 95% confidence limits, the one-sided lower limit on the number of participants with the ability is 22; we had set the true figure at 34. The one-sided lower limit on the long-term hit rate of the best participants is 0.51; our true figure was 0.75. Both of those limits seem reasonable to us.

For a real-data illustration, we use data from Wajid and Halpern (2012). Participants were asked which of 5 stimuli tasted different from the others when oleic acid had been added to just one of the 5. Each participant underwent two trials. Of 19 participants, 5 were correct on both trials, 9 on just one trial, and 5 on none. Testing the null of no discrimination, the upper-tail SILR test yielded $p_s = 0.00008$. (As an aside, TBB yielded 0.0004, some 5 times as large.) Dropping 4 of the best-performing participants one at a time raised p_s to 0.00062 then 0.0043 then 0.025 then 0.12. Because four of the best performers had to be removed to make the result non-significant at the 0.05 level, we conclude that 4 is the lower one-sided 95% limit on the number of the 19 participants who can detect oleic acid. Thus, there is no evidence that most people can detect

oleic acid in these circumstances, but the data suggest that a noticeable minority can do so. The lower 95% confidence limit on p_{max} is 0.353, so some people can apparently perform noticeably better than chance, though there is no convincing evidence that anyone can detect oleic acid on most trials.

Acknowledgements

Bruce P. Halpern called this problem to our attention, and made several valuable suggestions concerning usability. Harry T. Lawless offered several important suggestions concerning alternative methods of analysis. Douglas R. Elrod drew our attention to the convolution and FFT commands in R, and explained their use. All those suggestions were extremely helpful.

References

Beyene, Joseph (2001). Uses of the Fast Fourier Transform (FFT) in exact statistical inference.

PhD thesis, U. of Toronto. Online at

<https://tspace.library.utoronto.ca/bitstream/1807/16254/1/NQ63594.pdf>

Bi, Jian (2006). *Sensory discrimination tests and measurement*. Oxford, UK. Blackwell

Casella, George and Berger, Roger L. (2002). *Statistical inference*, 2nd ed. Duxbury/Thomson Learning, Pacific Grove, CA.

Cohen, Jacob (1977). *Statistical power analysis for the behavioral sciences*. New York: Academic Press.

Cooley, James W. and Tukey, John W. (1965). An algorithm for the machine calculation of complex Fourier series. *Mathematics of Computation*, 19, 297-301.

Edgington E.S. (2011) Randomization Tests. In: Lovric M. (eds) International Encyclopedia of Statistical Science. Springer, Berlin, Heidelberg

Edgington, Eugene S. and Onghena, Patrick(2007). Randomization tests, 4th edition. Boca Raton, FL: Chapman & Hall, CRC

Fisher RA (1925). Statistical methods for research workers. Edinburgh: Oliver and Boyd.

Fisher RA (1935). The design of experiments. Edinburgh: Oliver and Boyd.

Fisher RA (1936). "The coefficient of racial likeness" and the future of craniometry. *Journal of the Royal Anthropological Institute of Great Britain and Ireland*, 66:1, 57-63

Goldstine, Herman H. (1977). *History of Numerical Analysis from the 16th Through the 19th Century (Studies in the History of Mathematics and Physical Sciences, 2)*. New York: Springer-Verlag.

Lawless, Harry and Heymann, Hildegard (2010). *Sensory evaluation of food*. NY: Springer

Mood, Alexander M., Graybill, Franklin A., and Boes, Duane C. (1974). *Introduction to the theory of statistics*, 3rd ed. New York: McGraw-Hill.

Wajid, Naji A. and Halpern, Bruce P. (2012). Oral-cavity discrimination of vapor-phase long-chain 18-carbon fatty acids. *Chemical Senses*, 37:7, 595-602