PERMUTOOLS: A MATLAB PACKAGE FOR MULTIVARIATE PERMUTATION TESTING

Michael J. Crosse[1,2,3*]   **John J. Foxe**[3,4]   **Sophie Molholm**[3,4]

[1]Segotia, Galway, Ireland   [2]Trinity College Dublin, Ireland
[3]Albert Einstein College of Medicine, NY   [4]University of Rochester, NY
[*]crossemj@tcd.ie

ABSTRACT

Statistical hypothesis testing and effect size measurement are routine parts of quantitative research. Advancements in computer processing power have greatly improved the capability of statistical inference through the availability of resampling methods. However, many of the statistical practices used today are based on traditional, parametric methods that rely on assumptions about the underlying population. These assumptions may not always be valid, leading to inaccurate results and misleading interpretations. Permutation testing, on the other hand, generates the sampling distribution empirically by permuting the observed data, providing distribution-free hypothesis testing. Furthermore, this approach lends itself to a powerful method for multiple comparison correction — known as max correction — which is less prone to type II errors than conventional correction methods. Parametric methods have also traditionally been utilized for estimating the confidence interval of various test statistics and effect size measures. However, these too can be estimated empirically using permutation or bootstrapping techniques. Whilst resampling methods are generally considered preferable, many popular programming languages and statistical software packages lack efficient implementations. Here, we introduce PERMUTOOLS, a MATLAB package for multivariate permutation testing and effect size measurement.

Background

Hypothesis testing

For over a century, researchers have relied on parametric statistical procedures for conducting hypothesis testing, such as the famous Student's t-test [Student, 1908]. However, parametric testing was developed out of the necessity to make inferences about the null distribution, as it was impractical to generate it empirically. Today, it is possible to do so using permutation tests. Permutation tests work by permuting the observed data in an appropriate manner to compute the empirical distribution of the test statistic of interest — known as the permutation distribution — which approaches the null distribution [Fisher, 1935]. From the permutation distribution, we can estimate the confidence interval (CI) of the test statistic by computing the corresponding percentiles (e.g. 2.5% and 97.5% percentiles for 95% CI). We can also estimate the probability of observing such a result by chance (i.e. the p-value) by calculating the proportion of the permutation distribution that is greater than or equal to the magnitude of the test statistic (Fig. 1). The more permutations generated, the more accurate the estimate. Typically, permutations in the order of several thousand are required to obtain reliable p-values, which is computationally trivial for modern computers [Ernst, 2004].

Permutation testing can be applied to any statistical test. As no assumptions are made about the shape of the underlying distribution, permutation tests provide distribution-free nonparametric hypothesis testing without the need for rank-transformations [Holt and Sullivan, 2023]. For independent samples, permutation tests have been shown to be relatively insensitive to differences in population variance when samples of equal size are used [Murphy, 1967, Groppe et al., 2011b]. Moreover, permutation tests have been shown to be superior to parametric tests in biomedical research [Ludbrook and Dudley, 1998], and are more robust to the researcher's choice of significance level and thus more likely to produce replicable results [Noguchi et al., 2021]. Note, p-values should always be reported as a continuous measure (as opposed to the traditional dichotomous approach) but should never be interpreted as the strength

of an effect (for that, we use effect size measures) [Wasserstein and Lazar, 2016]. PERMUTOOLS offers permutation testing and confidence interval estimation for a range of statistical tests, including the ANOVA (one-way, two-way), *t*-test (one-sample, paired-sample, two-sample), *F*-test (two-sample), *Z*-test (one-sample), and correlation test (Pearson, Spearman, rankit). It does not output dichotomous test results to discourage this practice.

Correcting for multiple comparisons

When conducting multiple hypothesis tests simultaneously, there is an increased risk of false discoveries or type I errors. However, many of the traditional correction methods used to control family-wise error rate (FWER) tend to be overly conservative, resulting in increased type II errors (e.g. Bonferroni correction, Holm–Bonferroni method). Researchers are often faced with the delicate task of controlling the trade-off between type I and type II errors. In the biomedical field, researchers tend to preference controlling for type I errors, because the consequences of false positives can be detrimental, e.g. the introduction of an ineffective new therapy or treatment [Ludbrook and Dudley, 1998]. This inherent bias, along with the use of overly conservative correction methods, can be limiting in the pursuit of scientific progress.

Another advantage of permutation tests is that they can utilize a powerful technique for correcting for multiple comparisons — known as max correction — which is less prone to type II errors than conventional correction methods [Blair and Karniski, 1993, Westfall and Young, 1993]. Max correction (also known as t_{max} correction [Blair et al., 1994] or joint correction [Boca et al., 2014]) works as follows: on each permutation of the data, a separate test statistic is computed for each variable and the maximum absolute value (or most extreme positive or negative value) is taken across all variables. Repeating this procedure thousands of times produces a single permutation distribution against which the actual test statistic is compared (Fig. 1). Thus, taking the maximum across more variables naturally produces a more conservative permutation distribution. This highly intuitive approach provides strong control of FWER, even for small sample sizes [Gondan, 2010, Groppe et al., 2011a, Rousselet, 2023]. Max correction has been used in various scientific disciplines, including the study of electrophysiological data [Blair and Karniski, 1993, Groppe et al., 2011a,b] and human behavioural data [Gondan, 2010, Shaw et al., 2020, Crosse et al., 2022]. PERMUTOOLS automatically applies max correction to multivariate data, unless specified otherwise.

Measuring effect size

Effect size measurement is an equally important part of inferential statistics [Hentschke and Stüttgen, 2011]. Today, most scientific journals require that authors report the size of an effect, and not just its dichotomous existence. A common effect size measure used in research is the standardised mean difference, known as Cohen's *d* [Cohen, 1969]. Standardised effect sizes have the advantage of being metric-free, meaning that they can be directly compared across different studies. For independent samples, Cohen's *d* uses the pooled standard deviation when they are assumed to have equal variances, and an unpooled estimate when this cannot be assumed. For independent samples with significantly different variances, an estimate based on the control sample's variance can be used, known as Glass' delta [Glass, 1976]. In the case of ordinal data, an alternative formulation known as Cliff's *d* should be used [Cliff, 1993]. PERMUTOOLS gives the option to implement any of the above standardised effect size measures, as well as several unstandardised measures.

It is also typical to report the confidence interval of an effect size. Whilst they have traditionally been calculated using parametric methods, these too can be estimated empirically by generating the sampling distribution using a resampling procedure known as bootstrapping. Bootstrapping is fundamentally different to permutation testing in that it resamples with replacement, whereas permutation testing resamples without replacement. PERMUTOOLS uses an efficient bootstrapping algorithm that is optimised for multivariate data analysis.

Correcting for sample size

As with all statistical tests, the larger the sample size, the more accurately it can describe the population. However, in certain scientific disciplines, researchers are often required to work with relatively small ($n < 50$) sample sizes due to the limited availability of test subjects or other logistical study constraints. This can hinder our ability to measure the true, unbiased size of an effect. Despite the advantages of using a standardised effect size measure, metrics such as Cohen's *d* have been shown to have an upwards bias of up to about 4% for sample sizes of less than 50. This bias is somewhat reduced by using the pooled weighted standard deviation of the samples [Hedges, 1981]. Additionally, a bias correction factor can be applied to the effect size measure, which is approximately equal to $1 - 3/(4n - 9)$ [Hedges and Olkin, 1985]. When this correction factor is applied, it is usual to refer to the effect size as Hedges' *g*. The same correction factor can also be applied to the corresponding CIs. PERMUTOOLS automatically applies bias correction when calculating Cohen's *d* and Glass' delta (and their CIs), unless specified otherwise.

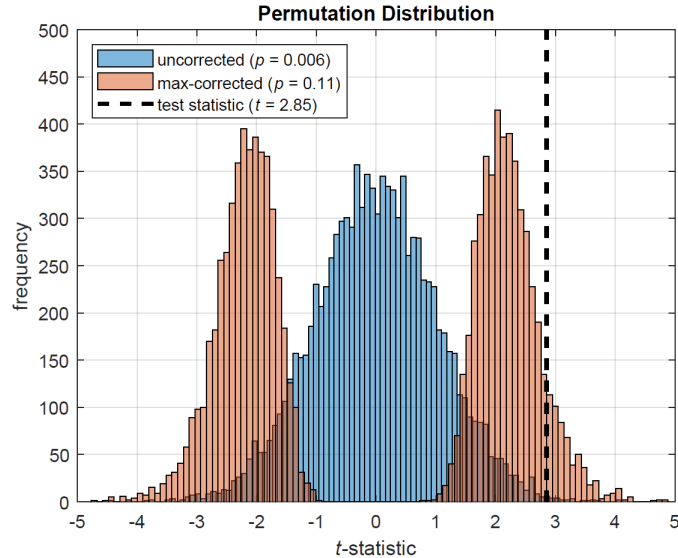**Permutation Distribution**

Figure 1: Permutation distributions for two-tailed tests based on the *t*-statistic (Blue: uncorrected, Red: max-corrected). Multivariate data were randomly generated to simulate two independent samples, each consisting of 20 variables (i.e. adjusted for 20 comparisons) and 30 observations (see example code below). The uncorrected permutation distribution shown is that for variable 3. The permutation test for variable 3 would not survive max correction at the 5% significance level ($p = 0.11$).

Statement of need

Whilst resampling methods have gained widespread acceptance, many popular programming languages, including MATLAB, lack efficient implementations or have not yet fully integrated them into their core statistical packages. This hinders the adoption of these robust and versatile techniques by researchers, therefore limiting the quality and reliability of quantitative research. To address this need, PERMUTOOLS provides a comprehensive set of functions in the MATLAB programming language for conducting resampling-based statistics that are easy to use and computationally efficient. Moreover, it is optimised for dealing with large, multivariate datasets and offers powerful methods for correcting for multiple comparisons and sample size, making it an invaluable tool for researchers across various fields of quantitative research. PERMUTOOLS offers a range of new features that distinguish it from existing statistical software packages. Some of the key features are described below.

Key features

- Optimised Resampling Algorithms: PERMUTOOLS utilizes efficient implementations of resampling algorithms that are optimised for multivariate data, ensuring efficient processing of even large datasets with negligible compute times.

- Multiple Comparison Correction: PERMUTOOLS implements the powerful max correction method to adjust *p*-values and CIs for multiple comparisons when dealing with multivariate data, reducing the risk of both type I and type II errors.

- Sample Size Correction: PERMUTOOLS applies a bias correction factor to the relevant standardised effect size measures and their CIs to adjust for any inherent inflation due to sample size, ensuring less biased estimates for small ($n < 50$) samples.

- Multivariate Processing: PERMUTOOLS handles multivariate datasets with ease, allowing for multiple tests to be performed simultaneously, as well as the option to perform pairwise comparisons between every combination of variables in a matrix (e.g. correlation matrix).

- Continuous Framework: PERMUTOOLS does not output test results under the traditional dichotomous decision-based framework (i.e. $H = 0, 1$). Instead, it outputs quantitative measures, encouraging researchers to interpret their results under a continuous framework.

Using PERMUTOOLS

The PERMUTOOLS functions are designed to mimic the API of the equivalent parametric functions in MATLAB, with the addition of the prefix 'permu' at the beginning of the function name. For example, to conduct a permutation test based on the *t*-statistic, the usual function `ttest()` becomes `permuttest()` etc. The input and output arguments are the same as before with the exception that the first output variable is always the test statistic (and not a dichotomous test result). An additional output variable containing the sampling distribution is included, as well as additional resampling-related input arguments that are described in the help documentation of each function. PERMUTOOLS uses a consistent input/output argument framework across all functions.

Example

The following example illustrates typical usage of the PERMUTOOLS toolbox. Here, we compare the means of two independent multivariate samples using permutation tests based on the *t*-statistic with max correction, and contrast the results with the equivalent parametric tests in MATLAB (i.e. two-sample *t*-tests). We then measure the associated effect sizes based on the bias-corrected standardised mean difference (i.e. Hedges' *g*), as well as the bootstrapped and parametric CIs.

First, we generate random multivariate data for two independent samples X and Y. Both samples have 20 variables, each with a mean value of approximately 0, except for the first 10 variables of Y which have a mean value of approximately -1. Each variable contains 30 observations.

```
% Generate random data
rng(42);
x = randn(30,20);
y = randn(30,20);
y(:,1:10) = y(:,1:10)-1;
```

Because we generated the random multivariate data from the same (normal) distribution, we can assume that their variances are approximately equal. If we could not assume this, we would first conduct a two-tailed test of variance based on the *F*-statistic using PERMUTOOLS' `permuvartest2()` function. Thus, we can proceed using the standard Student's *t*-statistic (as opposed to Welch's *t*-statistic). The two-sample permutation tests are implemented using PERMUTOOLS' `permuttest2()` function and the equivalent parametric tests are implemented using MATLAB's `ttest2()` function. By default, max correction is applied to the permutation tests.

```
% Run MATLAB's two-sample parametric t-test
[h1,p1,ci1,stats1] = ttest2(x,y);

% Run PERMUTOOLS' two-sample permutation t-test
[t2,p2,ci2,stats2] = permuttest2(x,y);
```

To measure the corresponding effect sizes, we compute Hedges' *g* and the 95% CIs using an efficient bootstrapping procedure. Effect size measures and bootstrapped CIs are calculated using PERMUTOOLS' `booteffectsize()` function. The equivalent parametric measures are calculated using MATLAB's `meanEffectSize()` function. By default, bias-correction is applied in both cases.

```
% Run MATLAB's parametric effect size measure
d3 = zeros(1,20);
ci3 = zeros(2,20);
for j = 1:20
    stats3 = meanEffectSize(x(:,j),y(:,j),'effect','cohen','paired',0);
    d3(j) = stats3.Effect;
    ci3(:,j) = stats3.ConfidenceIntervals';
end

% Run PERMUTOOLS' bootstrapped effect size measure
[d4,ci4,stats4] = booteffectsize(x,y,'effect','cohen','paired',0);
```

To compare the results of our parametric and permutation tests, we next plot some of the statistics as a function of variable (Fig. 2). The effect of max correction is clearly visible on the test statistic CIs and the resulting *p*-values, which are consistently more conservative than those of the uncorrected parametric tests (Fig. 2, left, middle). Importantly, comparisons where there appear to be spuriously significant effects in the parametric case (e.g. variable 20) did not survive the maximal statistic criterion in the permutation tests. In the effect size analysis, we can see that the bootstrapped CIs approximate the parametric CIs reasonably well (Fig. 2, right).

```matlab
% Set up figure
figure('Name','Permutation & Effect Size Analysis')
set(gcf,'color','w')
xaxis = 1:20;

% Plot test statistics & CIs
subplot(1,3,1), hold on
plot(xaxis,stats2.mu,'LineWidth',2.5)
plot(xaxis,ci1,'k',xaxis,ci2,'--r','LineWidth',1)
xlim([0,21]), ylim([-3,3]), box on, grid on
title('Test Statistic'), xlabel('variable'), ylabel('X-Y')
legend('mean difference','95% CI (param.)','','95% CI (perm.)')

% Plot p-values
subplot(1,3,2), hold on
plot(xaxis,p1,'k',xaxis,p2,'--r','LineWidth',1.5)
ylim([0,1]), xlim([0,21]), box on, grid on
title('{\itP}-value'), xlabel('variable'), ylabel('probability')
legend('{\itp}-value (param.)','{\itp}-value (perm.)')

% Plot effect sizes & CIs
subplot(1,3,3), hold on
plot(xaxis,d4,'LineWidth',2.5)
plot(xaxis,ci3,'k',xaxis,ci4,'--r','LineWidth',1)
ylim([-2,6]), xlim([0,21]), box on, grid on
title('Effect Size'), xlabel('variable'), ylabel('Hedges'' {\itg}')
legend('effect size','95% CI (param.)','','95% CI (boot.)')
```

From the above statistical analysis, we can examine the results of each individual pairwise comparison between X and Y, which are contained in the variables output by the PERMUTOOLS functions. Taking the first variables of X and Y as an example, we see that the mean of X_1 ($M = -0.06$, $SD = 0.91$) was significantly greater than the mean of Y_1 ($M = -1.09$, $SD = 0.86$), even after adjusting for multiple comparisons ($t(58) = 4.49$, $p = 0.0008$, Hedges' $g = 1.14$, 95CI [0.68, 1.72]). We recommend reporting the results of such statistical comparisons as shown here, in particular, with the effect size and CI included.

Future work

PERMUTOOLS is under active development, and new functions and features are constantly being added to it as needed. Future work aims to expand the toolbox to provide new permutation-based tests, including repeated measures ANOVAs (one-way, two-way) and multi-way ANOVAs (*n*-way), as well as expanding existing permutation tests to be able to deal with unbalanced and multivariate samples (e.g. unpaired ANOVAs). Whilst permutation tests have been shown to outperform nonparametric tests based on rank transformations [Holt and Sullivan, 2023], there are certain situations where a rank-transformation approach is desirable; for example, when dealing with ordinal data and outliers. To accommodate a wider range of data types, future work aims to develop permutation tests based on rank transformations, e.g. Sign test, Wilcoxon signed rank test, Mann-Whitney *U*-test, Kruskal-Wallis test, Friedman test, etc.

If you would like to contribute to PERMUTOOLS, or request that a specific statistical test or feature be added to it, please email the corresponding author at the email address provided above.
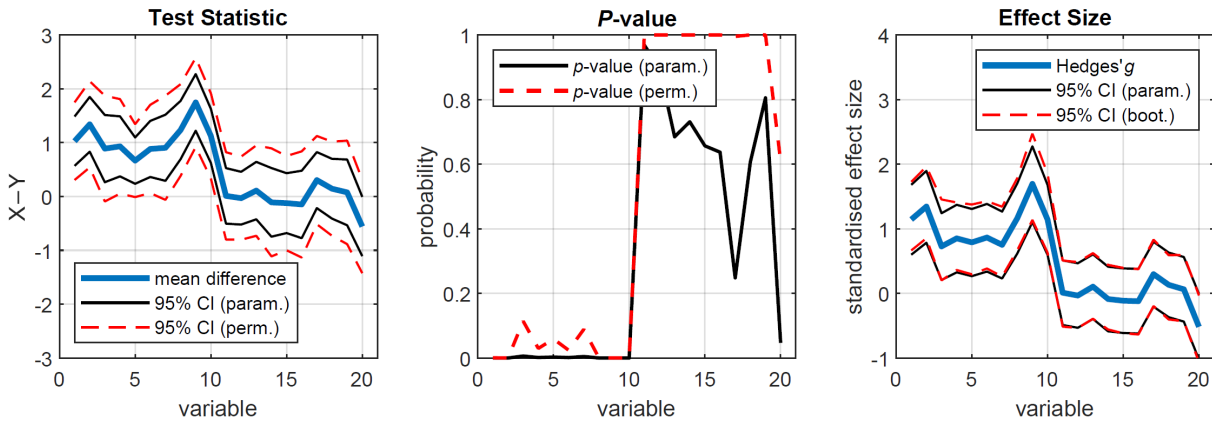
Figure 2: Results of the permutation and effect size analysis. Left: Test statistic (unstandardised) and 95% CI (parametric and permutation) for each test. Middle: *P*-value (parametric and permutation) for each test. Right: Effect size (Hedges' *g*) and 95% CI (parametric and bootstrapped) for each test.

Availability

PERMUTOOLS is freely available to download from GitHub (`https://github.com/mickcrosse/PERMUTOOLS`) and is compatible with Windows, Linux, and macOS. Example code demonstrating its usage is included in the examples folder and on the GitHub page.

Acknowledgments

Author contributions

M.J.C. conceived of the toolbox and manuscript, developed the toolbox and example code, and wrote the first draft of the manuscript. S.M. and J.J.F. supervised and edited the manuscript.

References

R Clifford Blair and Walt Karniski. An alternative method for significance testing of waveform difference potentials. *Psychophysiology*, 30(5):518–524, 1993.

R Clifford Blair, James J Higgins, Walt Karniski, and Jeffrey D Kromrey. A study of multivariate permutation tests which may replace hotelling's t2 test in prescribed circumstances. *Multivariate Behavioral Research*, 29(2):141–163, 1994. ISSN 0027-3171.

Simina M Boca, Rashmi Sinha, Amanda J Cross, Steven C Moore, and Joshua N Sampson. Testing multiple biological mediators simultaneously. *Bioinformatics*, 30(2):214–220, 2014.

Norman Cliff. Dominance statistics: Ordinal analyses to answer ordinal questions. *Psychological bulletin*, 114(3): 494, 1993.

Jacob Cohen. *Statistical power analysis for the behavioral sciences*. Academic press, 1969.

Michael J Crosse, John J Foxe, Katy Tarrit, Edward G Freedman, and Sophie Molholm. Resolution of impaired multisensory processing in autism and the cost of switching sensory modality. *Communications biology*, 5(1):601, 2022.

Michael D Ernst. Permutation methods: a basis for exact inference. *Statistical Science*, pages 676–685, 2004.

Ronald A Fisher. *The Design of Experiments*. Hafner, New York, 1935.

Gene V Glass. Primary, secondary, and meta-analysis of research. *Educational researcher*, 5(10):3–8, 1976.

Matthias Gondan. A permutation test for the race model inequality. *Behavior Research Methods*, 42(1):23–28, 2010. ISSN 1554-351X.

David M Groppe, Thomas P Urbach, and Marta Kutas. Mass univariate analysis of event-related brain potentials/fields i: A critical tutorial review. *Psychophysiology*, 48(12):1711–1725, 2011a. ISSN 0048-5772.

David M Groppe, Thomas P Urbach, and Marta Kutas. Mass univariate analysis of event-related brain potentials/fields ii: Simulation studies. *Psychophysiology*, 48(12):1726–1737, 2011b.

Larry V Hedges. Distribution theory for glass's estimator of effect size and related estimators. *journal of Educational Statistics*, 6(2):107–128, 1981.

Larry V Hedges and Ingram Olkin. *Statistical methods for meta-analysis*. Academic press, 1985.

Harald Hentschke and Maik C Stüttgen. Computation of measures of effect size for neuroscience data sets. *European Journal of Neuroscience*, 34(12):1887–1894, 2011.

Charles A Holt and Sean P Sullivan. Permutation tests for experimental data. *Experimental Economics*, pages 1–38, 2023.

John Ludbrook and Hugh Dudley. Why permutation tests are superior to t and f tests in biomedical research. *The American Statistician*, 52(2):127–132, 1998.

Brian P Murphy. Some two-sample tests when the variances are unequal: a simulation study. *Biometrika*, 54(3-4): 679–683, 1967.

Kimihiro Noguchi, Frank Konietschke, Fernando Marmolejo-Ramos, and Markus Pauly. Permutation tests are robust and powerful at 0.5% and 5% significance levels. *Behavior Research Methods*, 53(6):2712–2724, 2021.

Guillaume A Rousselet. Using cluster-based permutation tests to estimate meg/eeg onsets: how bad is it? *bioRxiv*, pages 2023–11, 2023.

Luke H Shaw, Edward G Freedman, Michael J Crosse, Eric Nicholas, Allen M Chen, Matthew S Braiman, Sophie Molholm, and John J Foxe. Operating in a multisensory context: Assessing the interplay between multisensory reaction time facilitation and inter-sensory task-switching effects. *Neuroscience*, 436:122–135, 2020.

Student. The probable error of a mean. *Biometrika*, 6(1):1–25, 1908.

Ronald L Wasserstein and Nicole A Lazar. The asa statement on p-values: context, process, and purpose. *The American Statistician*, 70(2):129–133, 2016.

Peter H Westfall and S Stanley Young. *Resampling-based multiple testing: Examples and methods for p-value adjustment*, volume 279. John Wiley & Sons, 1993.