Published in final edited form as:

Stat Med. 2017 July 20; 36(16): 2630–2640. doi:10.1002/sim.7288.

Tests for informative cluster size using a novel balanced bootstrap scheme

Jaakko Nevalainen^{a,*}, Hannu Oja^b, and Somnath Datta^c

^aSchool of Health Sciences, University of Tampere, Finland ^bDepartment of Mathematics and Statistics, University of Turku, Finland ^cDepartment of Biostatistics, University of Florida, United States

Abstract

Clustered data are often encountered in biomedical studies, and to date, a number of approaches have been proposed to analyze such data. However, the phenomenon of informative cluster size (ICS) is a challenging problem, and its presence has an impact on the choice of a correct analysis methodology. For example, Dutta and Datta (2015, *Biometrics*) presented a number of marginal distributions that could be tested. Depending on the nature and degree of informativeness of the cluster size, these marginal distributions may differ, as do the choices of the appropriate test. In particular, they applied their new test to a periodontal data set where the plausibility of the informativeness was mentioned but no formal test for the same was conducted. We propose bootstrap tests for testing the presence of ICS. A balanced bootstrap method is developed to successfully estimate the null distribution by merging the re-sampled observations with closely matching counterparts. Relying on the assumption of exchangeability within clusters, the proposed procedure performs well in simulations even with a small number of clusters, at different distributions and against different alternative hypotheses, thus making it an omnibus test. We also explain how to extend the ICS test to a regression setting and thereby enhancing its practical utility. The methodologies are illustrated using the periodontal data set mentioned above.

Keywords

Bootstrapping; Hypothesis testing; Clustered data; Informative cluster size; Matching

1. Introduction

Clustered data frequently arise in biomedical research. Repeated measurements on the same individual, patients at different hospitals or animals from a number of litters are examples where the observations are dependent within a cluster, but the clusters can be assumed to be independent. Methods for modeling such data are well developed and recognized [1, 2, 3]. General methodology assumes that the cluster size is a fixed design constant. However, it has been noted, for example, in volume-outcome studies that specialized surgeons treating

*Correspondence to: School of Health Sciences, University of Tampere, Finland. jaakko.nevalainen@uta.fi.

many patients may have better outcomes than those treating few patients [4]; in periodontal studies that patients with fewer teeth tend to have poorer condition of the teeth that remain [5, 6]; and in animal studies that treatments may have an effect on fetal weight with or without an effect on fetal losses [7].

It is important to appreciate that the cluster size could be informative for inference. First, a proper analysis of informative cluster size (ICS) data involves the consistent estimation of quantities of interest for the desired target population. The latter can be defined as a typical member of a typical cluster in the population or a typical member of the population [8]. With missing data, the target of inference could be on hypothetical clusters that are observed incompletely by contrasting complete cluster inference with observed cluster inference [9]. In general, the parameters of interest are not the same for alternative target populations with ICS, and their estimation thus requires correctly specified estimators [9, 10]. Second, standard methods for clustered data analysis may result in incorrect sizes and biased estimates with ICS, and conversely, ICS methods can lead to a loss of efficiency if the cluster size is non-informative [11]. Therefore, it is crucial to both detect ICS and, in presence of ICS, define the target population of inference. As for example, Dutta and Datta [12] recently presented a number of marginal distributions that could be tested. Depending on the nature and degree of informativeness of the cluster size, these marginal distributions may differ, as do the choices of appropriate rank test. In particular, they applied their new test to a periodontal data set, where the plausibility of the informativeness was mentioned but no formal test for the same was conducted. In Section 6 of this paper, we use this data set and show that the cluster size is indeed informative by applying our ICS test.

It is possible to investigate the informativeness of the cluster size descriptively by a comparison of the marginal distributions between strata defined by cluster size [13]. We address the ICS detection problem more formally and propose an omnibus test for informativeness of the cluster size. Benhin et al. [11] investigated a Wald-type test for ignorability of cluster size in a particular setting, the estimating equations framework for linear and logistic regression models. The size of their tests were close to the nominal level. To the best of our knowledge, no other test for the presence of ICS problems currently exists. In general, the null distributions of the test statistics for ICS are analytically intractable, and novel bootstrap strategies are needed because the standard bootstrap turns out to be problematic. Simple solutions, such as using the cluster size as a covariate in the mean specification of the model, are approximate at best. The simulation models considered below are examples of such models. We demonstrate the validity and efficiency of the proposed tests for cases both without and with covariates.

The paper is organized as follows. Section 2 defines the problem and gives possible test statistics. The balanced bootstrap procedure is proposed in Section 3, with an extension to regression in Section 4. Simulation results and applications to dental data are presented in Sections 5 and 6, respectively. Section 7 completes the paper with a discussion. Full details of the balanced bootstrap algorithm are placed in the Appendix.

2. Setup and test statistics

Let $(V_1, \dots, V_M)$ be a sample of independent and identically distributed (iid) observations from an unknown distribution, where each V_j represents a cluster consisting of

$$(N_i; (Y_{i1}, X_{i1}), \dots, (Y_{iN_i}, X_{iN_i})),$$

the cluster size N and the outcomes Y plus covariates X . Within a cluster, the outcomes, given the covariates, are dependent, and their marginal distribution is potentially dependent on the cluster size.

In the case of no covariates, the interest is on the properties of the marginal distribution of the outcome. Informative cluster size refers to any violation of the condition

$$P(Y_{ij} \leq y | N_i = k) = P(Y_{ij} \leq y), \quad k=1, 2, \dots, K; j=1, 2, \dots, k,$$

which would mean that the marginal distributions depend on the cluster size. Nevalainen et al. [10] considered the tools for inference on the marginal distribution with ICS. If the condition holds, the cluster size is irrelevant, and a standard analysis of clustered data can be used.

With covariates, cluster size is noninformative if the marginal distribution conditional on the covariates is not influenced by the cluster size [13]:

$$P(Y_{ij} \leq y | X_{ij}, N_i = k) = P(Y_{ij} \leq y | X_{ij}), \quad k=1, 2, \dots; j=1, \dots, k, \quad (1)$$

If there are no covariates, the null hypothesis that is implied by within-clusters exchangeability and to be tested is simply

$$H_0: P(Y_{ij} \leq y | N_i = k) = F(y) \text{ for all } j \leq k, k=1, \dots, K, \text{ for some unknown distribution } F,$$

which means that under the null hypothesis, all of the marginal distributions are the same and do not depend on N_i . Note, however, that the stronger null hypothesis

$$H_0^*: P(Y_{i1} \leq y_1, \dots, Y_{ij} \leq y_j | N_i = k) = F(y_1, \dots, y_j), \text{ for all } j \leq k, k=1, \dots, K,$$

may not be true. Although this null hypothesis of same joint (sub-)distributions may be more relevant if, for example, intracluster correlation varies across elements or if the cluster members are ordered within a cluster (such as in time), we do not pursue this direction in this paper.

Test statistics for testing differences in the conditional distributions, given cluster sizes, include standard tests for multiple samples, treating cluster size as the grouping factor. The F -test derived from the one-way analysis of variance and the nonparametric Kruskal-

Wallis test statistics are used as benchmark tests. Their weakness is that these tests were designed to detect shifts in location.

An omnibus test would have power against a wider range of alternatives such as differences in scale or other parameters characterizing the distribution. Therefore, it is reasonable to look for any differences in the conditional distributions of Y given N_i .

There are many possibilities for constructing omnibus test statistics as multisample versions of Kolmogorov-Smirnov or Cramér von Mises multisample test-type test statistics [14]. In terms of broad validity and power across several investigations, we suggest the use of

$$T_F = \sup_y |\hat{F}(y) - \tilde{F}(y)|,$$

where the “null estimator” of the distribution function

$$\hat{F}(y) = \left(\sum_{i=1}^M N_i \right)^{-1} \sum_{i=1}^M \sum_{j=1}^{N_i} (Y_{ij} \leq y)$$

is contrasted with an alternative estimator

$$\tilde{F}(y) = M^{-1} \sum_{i=1}^M N_i^{-1} \sum_{j=1}^{N_i} (Y_{ij} \leq y).$$

Under the null hypothesis, these two estimators are consistent estimators of the same population distribution function. As a competitor to T_F we propose a version of the Cramér von Mises multisample test:

$$T_{CM} = \sum_{k \in \mathcal{J}} \left[k M_k \int \left(\hat{F}_k(y) - \hat{F}(y) \right)^2 dy \right],$$

where $\mathcal{J}$ is the set of unique cluster sizes, M_k is the number of clusters of size k , and

$$\hat{F}_k(y) = \frac{1}{k M_k} \sum_{i=1}^M \sum_{j=1}^{N_i} (N_i = k, Y_{ij} \leq y)$$

is the estimator of the distribution for cluster size k . This test statistic is potentially more powerful than T_F . The same test statistics can be used in the regression setting to test the presence of ICS as described in Section 4.

3. Balanced bootstrap procedure

The distribution of the test statistics under the null is in general analytically intractable. Therefore, the use of the bootstrap to approximate the null distribution is a possibility. However, even the null bootstrap for clustered data with ICS turns out to be challenging.

Because the V_i 's are iid, the first idea could be to bootstrap them directly [15]. One could have either (i) bootstrap samples of size M from $V_1, \dots, V_M$ such that $M_1, M_2, \dots$ are random or (ii) separate bootstrap samples of size M_k from a subset of $V_1, \dots, V_M$ with cluster size k , $k = 1, 2, \dots$. In practice, however, these strategies do not succeed: they are limited to settings in which the number of clusters is large relative to the number of distinct values of N_i 's in the sample. Otherwise, the distribution based on the bootstrap sample lacks sufficient richness. At one extreme, consider a data set with one cluster of size k , and all $M - 1$ other clusters of size k' . For any test statistic in the previous section, the contribution of cluster size k to the test statistic would be identical in all bootstrap samples, and the procedure would not replicate variation under the null for that particular cluster size. A proper bootstrap procedure for clustered data must be able to address many distinct values of N_i relative to the number of clusters.

A successful null bootstrap can be constructed as follows. Because the interest is only in the conditional distributions of $Y_{i1}, \dots, Y_{iN_i}$ given N_i , then conditioning on ancillary statistics $N_1, \dots, N_M$ is a compelling approach. We therefore propose the following conditional bootstrap testing scheme:

1. Choose the test statistic $T = T(V_1, \dots, V_M)$, where $(V_1, \dots, V_M)$ for testing ICS.
2. Re-sample B bootstrap samples by repeating for every $j = 1, \dots, B$:
 - Permute the data set by permuting cluster members within each cluster.
 - Re-sample M clusters from the permuted data set by repeating for every $i = 1, \dots, M$:
 - Draw a random cluster V with index i^* from $\{1, \dots, M\}$.
 - If $N_{i^*} = N_i$ then the bootstrap cluster is $V_{ji}^* = (N_i; Y_{i*1}, \dots, Y_{i*N_i})$.
 - If $N_{i^*} < N_i$ then the bootstrap cluster is merged from two “matching” clusters:

$$V_{ji}^* = (N_i; Y_{i*1}, \dots, Y_{i*N_{i^*}}, Y_{k(N_{i^*}+1)}, \dots, Y_{kN_i}),$$

$$\text{where } k = \arg \min_k \{D(V_{i^*}, V_k) : N_k = N_i\}.$$

- The j th bootstrap sample is $V_j^* = (V_{j1}^*, \dots, V_{jM}^*)$.
- Compute the value of the test statistic for the bootstrap sample $T_j^* = T(V_j^*)$.

3. The null distribution of the test statistic can be approximated by the obtained $T_1^*, \dots, T_B^*$. Finally, the p-value is computed as

$$\frac{1}{B} \sum_{j=1}^B (T_j^* \geq T).$$

The permutation in step 2 ensures a rich structure by generating bootstrap samples j^* from cluster-wise permuted observations, $j = 1, \dots, B$. With this step, initial clusters can be matched with different counterparts in different bootstrap samples, for example. By merging the matching clusters, the cluster sizes of V_{ji}^* are still N_i for all $j = 1, \dots, B$ (obeying the conditionality principle). Matching also addresses the dependence between $Y_{i1}, \dots, Y_{iN_i}$ because exchangeability does not imply independence. The distance between two clusters used to identify counterparts is simply defined as

$$D(V_i, V_j) = (\min\{N_i, N_j\})^{-1} \sum_{k=1}^{\min\{N_i, N_j\}} (Y_{ik} - Y_{jk})^2.$$

In case of tied distances (with aggregated data, for instance), choose one of the clusters to minimize the distance at random. The algorithm is given in explicit detail in the Appendix.

The procedure requires re-calculation of distances at each bootstrap sample and is thus computationally intensive. Whether this is an issue depends mainly on the sample size and may be controlled to some extent by keeping B reasonably small such as 500. We have written an R [16] function for conducting the test for ICS, and the code is available as supplementary online material.

4. Extensions to the regression setting

First, we consider an additive error regression setting

$$h(Y_{ij}) = g(\theta; X_{ij}) + \varepsilon_{ij}; 1 \leq j \leq N_i, i \geq 1,$$

where $\theta \in \mathbb{R}^p$, for some $p \geq 1$, is a vector of regression parameters, h and g are known real-valued transformation and regression functions and ε are additive errors, independent of the X .

To understand the potential informativeness of the cluster size and the construction of the model residuals under ICS, it is useful to consider a mixed effects model, where we assume that the errors ε_{ij} comprise cluster-level unobserved random effects μ_i and subject-level iid random errors η_{ij} , i.e., $\varepsilon_{ij} = \mu_i + \eta_{ij}$, $0 = E(\mu_i) = E(\eta_{ij})$. A violation of the ICS condition (1) will occur, for example, if the cluster sizes and μ_i 's are dependent.

We formulate a model-based re-sampling procedure that can be used to extend our previous tests, as developed in Sections 3 and 4, to test ICS for the marginal distribution of the

response Y to this regression setting. The procedure can be algorithmically described as follows:

1. Estimate the regression parameters using ordinary least squares, leading to the estimates of the regression parameters and the corresponding model residuals:

$$\hat{\theta} = \arg \min_{\theta} \sum_{i=1}^M \sum_{j=1}^{N_i} \{Y_{ij} - g(\theta, X_{ij})\}^2, \\ e_{ij} = Y_{ij} - g(\hat{\theta}, X_{ij}); 1 \leq j \leq N_i, 1 \leq i \leq M.$$

2. Compute the value of a given test statistic T as described in Section 3, except we replace the data Y_{ij} with the model residuals e_{ij} .
3. Resample the clusters of data as above, leading to bootstrap data

$$V_i^* = (N_i^*; (Y_{ij}^*, X_{ij}^*), 1 \leq j \leq N_i^*).$$

4. Recompute the estimate $\hat{\theta}$ and the residuals using the bootstrap data obtained in step 3.

$$\hat{\theta}^* = \arg \min_{\theta} \sum_{i=1}^M \sum_{j=1}^{N_i^*} \{Y_{ij}^* - g(\theta, X_{ij}^*)\}^2, \\ e_{ij}^* = Y_{ij}^* - g(\hat{\theta}^*, X_{ij}^*); 1 \leq j \leq N_i^*, 1 \leq i \leq M.$$

and compute the test statistic T , denoted T_j^* , for the j th bootstrap sample using these residuals $\{e_{ij}^*\}$

5. Compute the null p-value from independent Monte Carlo replications of the T_j^* , $j = 1, \dots, B$.

An alternative test for ICS can be constructed by using the quadratic form of the difference between the unweighted regression parameter estimate and the ICS-resistant estimate of the regression parameter.

5. Simulations

Next, we investigate how well the tests hold their nominal levels and compare their efficiencies. The simulations at the null cases involve $M = 25, 50$ and 100 clusters, 500 bootstrap samples and 5000 Monte Carlo samples. For efficiency comparisons, we used 500 or 1000 Monte Carlo runs, depending on the computational expense of the simulation in question. Simulations were carried out on a supercluster provided by CSC – IT Center for Science Ltd., Finland. The null distributions of all test statistics were obtained by approximating them using the proposed bootstrap procedure. Simulation model A allows infinitely many cluster sizes, whereas simulation model B is limited to a fixed number of possible cluster sizes.

Simulation model A: infinitely many cluster sizes

Simulation model A represents a case in which the researcher cannot control the cluster sizes, which are assumed to arise from a distribution. This is a realistic scheme in observational studies in particular.

Data are generated as follows. Under the null hypothesis, set $\mu_i = 0$; in the alternative case, generate a cluster-specific latent variable $\mu_i \stackrel{iid}{\sim} N(0, 1), i=1, \dots, M$ that is responsible for ICS. Then, V_j entities are simulated by

$$\begin{aligned} \lambda_i &= \lambda \exp(\gamma \mu_i) \\ (N_i - 1) &\stackrel{iid}{\sim} \text{Poisson}(\lambda_i), i=1, \dots, M \\ Y_{ij} &\stackrel{iid}{\sim} N(\mu_i + \beta X_{ij}, 1), j=1, \dots, N_i \end{aligned}$$

In this model, a setting without covariates is obtained with $X_{ij} = 1$ for all i, j , and with covariates, the X_{ij} are set as iid from a Bernoulli(1/2) distribution.

Simulation model B: two possible cluster sizes

This model demonstrates the case where cluster sizes are fixed to two possible values. Settings such as these are encountered in designed experiments and survey samples. First, draw $\mu_i \stackrel{iid}{\sim} N(0, 1)$ and $\alpha_i \stackrel{iid}{\sim} N(0, 1)$. Clusters $V_i, i=1, \dots, M$, are then generated as

$$\begin{aligned} N_i &= \begin{cases} 5, & \text{if } \mu_i > a; \\ 10, & \text{if } \mu_i < a. \end{cases} \\ Y_{ij} &= \begin{cases} \alpha_i + \beta X_{ij} + \varepsilon_{ij}, & \text{if } N_i = 5; \\ \alpha_i + (\beta + \gamma) X_{ij} + \zeta_{ij}, & \text{if } N_i = 10, \end{cases} \end{aligned}$$

where ε_{ij} are iid from t_{v_1} and ζ_{ij} are iid from t_{v_2} , thus making it possible to use alternative error distributions. The parameter a controls the balance in terms of the percentage of clusters that were of size 5 (10). The setting without covariates is with $X_{ij} = 1$ for all i, j , and with covariates, the X_{ij} are iid from Bernoulli(1/2). Under the null hypothesis, $\gamma = 0$ and $v_1 = v_2$.

5.1. Results

Among the test statistics considered, T_F proved to be the most conservative (Table 1). It tended to be conservative in settings both with and without covariates and in all of the simulations. Changes in the number of clusters did not bring it closer to the nominal level. The other tests, however, were all liberal in Simulation A. The number of different cluster sizes increases rapidly as a function of the λ -parameter, and all three tests appear to suffer from this. A larger M even makes this worse. The reason is that unlike T_F these three tests are based on grouping by cluster size. Some groups are unavoidably small in this particular simulation. Simulation B consists of only two cluster sizes. Here the three tests are very close to the nominal level, and there seems to be a small improvement as the number of clusters increases. The results obtained for regression settings do not differ from those

obtained without covariates. Based on Simulation B, an imbalance in the number of clusters of different size does not seem to result in a worse performance.

We performed additional simulations to investigate the behavior of the test in presence of autocorrelated errors (provided as supplementary online material). These simulations suggest that the levels of the tests remain similar to those reported in Table 1, and that the balanced bootstrap test is not vulnerable to a modest departure from the exchangeability assumption.

Figure 1 shows the empirical power of two tests against the sequence of alternatives from $\gamma = 0.05$ to 0.5. Against this type of alternative, T_F is more powerful than T_{CM} , which is not surprising because of the sparsity of the distribution of cluster sizes in Simulation A. Already with $M = 25$ clusters, the tests demonstrate some power, and with 100 clusters, the deviations from the null are detected with high probabilities. Power results for the F -test and Kruskal-Wallis are not shown because of their poor validity under the null hypothesis (Table 1).

Table 2 reports the empirical powers of the tests for Simulation B. Parameters ν_1 and ν_2 determine how heavy the tails of the t -distribution are; extreme choices of $\nu_2 = 1$ and $\nu_1 = \infty$ yield the Cauchy and normal distributions, respectively. None of the tests has sufficient power to detect the differences for any choice of ν_2 with 25 clusters. At $M = 100$, two test statistics exhibit power against the tail alternative; among them, T_{CM} is superior to T_F . Recall that T_{CM} was obtained by integrating the area between the empirical distribution functions over the whole range of data, whereas T_F is based on the maximal difference between the estimates of the distribution function. This explains the superior performance: with the tail alternative, the difference exists for the whole data range, not just locally. The F -test performs oddly at $\nu_2 = 1$ with no power. The reason for the notoriously poor performance is that the moments for the Cauchy distribution do not exist. The Kruskal-Wallis test avoids this problem because it operates with ranks but has no power. This is expected because the alternative is not a location shift.

Already at $M = 25$, all tests show power against the regression alternative controlled by parameter γ . There are no substantial differences between the test statistics, but T_F appears to be slightly less powerful than the others.

6. Numerical illustrations using dental data

We illustrate our testing methodology by using a periodontal data set extracted from the Piedmont 65+ Dental Study [17]. The Piedmont Health Study of the Elderly [18], which is the parent study for this Piedmont 65+ Dental Study, is a longitudinal study of the health status of people aged 65 and over in five contiguous North Carolina counties. The Piedmont 65+ Dental Study takes advantage of the data available from the parent study and collects additional information by means of an interview, oral examination, and microbiological and salivary assays.

For an illustration of our tests for ICS without covariates, we consider the data collected at baseline. Our response variable is the total attachment loss, which is measured at the tooth

level. Attachment loss is one measure of the severity of destructive periodontal disease in a patient. It is measured as the distance in millimeters from the cemento-enamel junction to the base of the pocket (represented by the bottom of the probe during a dental examination). The analysis was performed on a subsample based on 100 randomly selected patients (clusters), each having from 1 to 32 teeth (median number of teeth: 19). The patients' mean attachment loss (across teeth in a patient's mouth) is plotted in Figure 2. It is apparent that the mean attachment loss tends to be greater for those patients with less teeth (smaller cluster size); that is, the cluster size is potentially informative.

Our bootstrap test verifies this finding: the T_F and T_{CM} tests reject the ICS null hypothesis, with a p-value less than 0.01 (based on 1000 bootstrap samples). It took 19 minutes and 1 second to carry the test out on a PC (Intel Core i7-6820 CPU 2.70GHz processor, 16 GB of RAM).

Therefore, any subsequent analysis on attachment loss should consider informative cluster size. Figure 3 shows the gap between the conventional empirical distribution function and a more reasonable descriptive quantity: the empirical distribution of random teeth in a random cluster. The latter was obtained by an inverse-cluster-size weighted estimate of the cumulative distribution function [10].

To illustrate a regression setting, two regression models were fitted, using the square root of the attachment loss as the response and smoking status (no/yes) and socioeconomic index (SEI) as explanatory variables. The first fit was an ordinary least squares (OLS) fit, and the second one used the inverses of the cluster sizes as weights (WLS). Robust standard errors were computed for the estimates of the regression coefficients with patients as clusters. The estimates of the intercept and the regression coefficients are mostly similar, but the estimate for SEI seems slightly different in absolute value (Table 3). The balanced bootstrap test indicates the presence of ICS: $p = 0.001$ with T_{CM} and $p < 0.001$ with T_F . The one based on weighted estimates would therefore be the preferred result because the results based on ICS are potentially biased for an average tooth of a typical patient.

7. Discussion

In this paper, we have shown that the balanced bootstrap procedure approximates the null distribution of the test statistics in the ICS setting reasonably well. The procedure turned out to be liberal for three of the four test statistics that were considered in settings where the number of distinct cluster sizes was large relative to the number of clusters. In those cases, we would recommend using a grouping of cluster sizes to adjacent cluster sizes that were encountered, say, fewer than three times in the sample. One of the test statistics, T_F , was very close to or below its nominal level in that setting.

Among the test statistics considered, the multisample Cramér von Mises held its level fairly well and was powerful in all of the simulations considered. For a small number of groupings by cluster size, it would be the preferred test, but it was liberal when the cluster sizes had a large number of possible values. If the number of groupings is large and the number of clusters of each size is small, T_F is more useful for the following reasons: (i) it is

conservative rather than liberal, and (ii) it can be modified for specific alternatives easily. For example, it is unimportant to know whether ICS is present in the sense of shift, if the quantity of interest (e.g., correlation, regression coefficient) is not affected by that shift. T_F -type statistics can be used to test whether the parameter of interest suffers from ICS. Standard test statistics such as those conventionally used in the F -test and the Kruskal-Wallis test are limited to shift alternatives and performed very poorly in one of the simulations. This was observed even though their null distributions were based on the balanced bootstrap procedure, rather than the ones routinely used in one-way analysis of variance.

Supplementary Material

Refer to Web version on PubMed Central for supplementary material.

Acknowledgments

This research was supported by NIH grants 1R03DE020839 and 1R03DE022538. The authors would like to thank Jim Beck and Kevin Moss in the School of Dentistry at the University of North Carolina for providing the data set on periodontal disease from the Piedmont 65+ Dental study.

Appendix

Conditional bootstrap without covariates

1. Find $T=T()$, where $=(V_1, \dots, V_M)$. Let

$$V_i(\alpha_i) = (N_i; Y_{i\alpha_{i1}}, \dots, Y_{i\alpha_{iN_i}}),$$

where $\alpha_i = (\alpha_{i1}, \dots, \alpha_{iN_i})$ is a permutation of $(1, \dots, N_i)$, denote a permuted cluster and

$$(\alpha) = (V_1(\alpha_1), \dots, V_M(\alpha_M))$$

the whole permuted data within clusters.

2. For $j = 1, \dots, B$, do
 - 2.1 $\leftarrow (\alpha)$ with random α .
 - 2.2 For all $i = 1, \dots, M$, do
 - 2.2.1 Choose a random V with index i^* .
 - 2.2.2 If $N_{i^*} = N_i$ then $V_{ji}^* = (N_i; Y_{i^*1}, \dots, Y_{i^*N_i})$.
 - 2.2.3 If $N_{i^*} < N_i$ then $V_{ji}^* = (N_i; Y_{i^*1}, \dots, Y_{i^*N_{i^*}}, Y_{k(N_{i^*}+1)}, \dots, Y_{kN_i})$ where $k = \arg \min_k \{D(V_{i^*}, V_k): N_k = N_i\}$.
 - 2.3 Obtain $_j^* = (V_{j1}^*, \dots, V_{jM}^*)$.
 - 2.4 Compute $T_j^* = T(_j^*)$.

3. Obtain $T_1^*, \dots, T_B^*$.
4. Find the p-value as $\frac{1}{B} \sum_{j=1}^B (T \geq T_j^*)$.

Conditional residual bootstrap in regression

Write

$$=(U_1, \dots, U_M) \text{ where } U_i = (N_i; (Y_{i1}, X_{i1}), \dots, (Y_{iN_i}, X_{iN_i})), i=1, \dots, M$$

and

$$=(V_1, \dots, V_M) \text{ where } V_i = (N_i; \varepsilon_{i1}, \dots, \varepsilon_{iN_i}), i=1, \dots, M$$

with the estimated residuals

$$\varepsilon_{ij} = Y_{ij} - g(\hat{\theta}, X_{ij}), i=1, \dots, M; j=1, \dots, N_i.$$

Here $\hat{\theta} = \hat{\theta}()$ is invariant under permutations within clusters.

As before, write

$$U_i(\alpha_i) = (N_i; (Y_{i\alpha_{i1}}, X_{i\alpha_{i1}}), \dots, (Y_{i\alpha_{iN_i}}, X_{i\alpha_{iN_i}})) \text{ and } V_i(\alpha_i) = (N_i; \varepsilon_{i\alpha_{i1}}, \dots, \varepsilon_{i\alpha_{iN_i}})$$

where $\alpha_i = (\alpha_{i1}, \dots, \alpha_{iN_i})$ is a permutation of $(1, \dots, N_i)$, and finally

$$(\alpha) = (U_1(\alpha_1), \dots, U_M(\alpha_M)) \text{ and } (a) = (V_1(\alpha_1), \dots, V_M(\alpha_M))$$

for all permutations $a = (a_1, \dots, a_M)$.

Algorithm

-
- 1 For observed, find the residuals and $T = T()$.
 - 2 For $j = 1, \dots, B$, do
 - 2.1 $\leftarrow (\alpha)$ and $\leftarrow (\alpha)$ with random α .
 - 2.2 For all $i = 1, \dots, M$, do
 - 2.2.1 Choose a random V with index i^* .
 - 2.2.2 If $N_{i^*} \geq N_i$ then $U_{ji}^* = (N_i; (Y_{i^*1}, X_{i^*1}), \dots, (Y_{i^*N_i}, X_{i^*N_i}))$.
 - 2.2.3 If $N_{i^*} < N_i$ then
 - 2.2.3.1 $U_{ji}^* = (N_i; (Y_{i^*1}, X_{i^*1}), \dots, (Y_{i^*N_i}, X_{i^*N_i}), (Y_{k(N_{i^*}+1)}, X_{k(N_{i^*}+1)}), \dots, (Y_{kN_i}, X_{kN_i})),$
 - 2.2.3.2 where $k = \arg \min_k \{D(V_{i^*}, V_k) : N_k \geq N_i\}$.

- 2.3 Obtain $j^* = (U_{j1}^*, \dots, U_{jM}^*)$.
- 2.4 For j^* , find the residuals j^* and $T_j^* = T(j^*)$.
- 3 Obtain $T_1^*, \dots, T_B^*$.
- 4 Find the p-value as $\frac{1}{B} \sum_{j=1}^B (T \geq T_j^*)$.

References

1. Diggle, P., Heagerty, P., Liang, KY., Zeger, S. Analysis of Longitudinal Data. Second. Oxford University Press; 2002.
2. Fitzmaurice, G., Laird, N., Ware, J. Applied Longitudinal Analysis. John Wiley & Sons; 2004.
3. Aerts, M., Geys, H., Molenberghs, G., Ryan, LM., editors. Topics in Modelling of Clustered Data. Chapman & Hall/CRC; Boca Raton, Florida, USA: 2002.
4. Panageas KS, Schrag D, Russell LA, Venkatraman ES, Begg CB. Properties of analysis methods that account for clustering in volume-outcome studies when the primary predictor is cluster size. *Statistics in Medicine*. 2007; 26:2017–2035. [PubMed: 17016864]
5. Williamson JM, Datta S, Satten GA. Marginal analyses of clustered data when cluster size is informative. *Biometrics*. 2003; 59:36–42. [PubMed: 12762439]
6. Wang M, Kong MK, Datta S. Inference for marginal linear models for clustered longitudinal data with potentially informative cluster sizes. *Statistical Methods in Medical Research*. 2011; 20:347–367. DOI: 10.1177/0962280209347043 [PubMed: 20223781]
7. Dunson DB, Chen Z, Harry J. A Bayesian approach for joint modeling of cluster size and subunit-specific outcomes. *Biometrics*. 2003; 59:521–530. [PubMed: 14601753]
8. Huang Y, Leroux B. Informative cluster sizes for subcluster-level covariates and weighted generalized estimating equations. *Biometrics*. 2011; 67(3):843–851. [PubMed: 21281273]
9. Seaman SR, Pavlou M, Copas AJ. Methods for observed-cluster inference when cluster size is informative: A review and clarifications. *Biometrics*. 2014; 70:449–456. [PubMed: 24479899]
10. Nevalainen J, Datta S, Oja H. Inference on the marginal distribution of clustered data with informative cluster size. *Statistical Papers*. 2014; 55:71–92. [PubMed: 25878396]
11. Benhin E, Rao JNK, Scott AJ. Mean estimating equation approach to analysing cluster-correlated data with nonignorable cluster sizes. *Biometrika*. 2005; 92:435–450. DOI: 10.1093/biomet/92.2.435
12. Dutta S, Datta S. A rank-sum test for clustered data when the number of subjects in a group within a cluster is informative. *Biometrics*. 2016; 72:432–440. DOI: 10.1111/biom.12447 [PubMed: 26575695]
13. Hoffman EB, Sen PK, Weinberg CR. Within-cluster resampling. *Biometrika*. 2001; 88:1121–1134.
14. Kiefer J. K-sample analogues of the Kolmogorov-Smirnov and Cramér v. Mises tests. *The Annals of Mathematical Statistics*. 1959; 30:420–447.
15. Field CA, Welsh AH. Bootstrapping clustered data. *Journal of the Royal Statistical Society, Series B*. 2007; 69:369–390.
16. R Core Team. R: A Language and Environment for Statistical Computing. R Foundation for Statistical Computing; Vienna, Austria: 2016. URL <https://www.R-project.org/>
17. Beck JD, Koch GG, Rozier RG, Tudor GE. Prevalence and risk indicators for periodontal attachment loss in a population of older community-dwelling blacks and whites. *Journal of Periodontology*. 1990; 61:521–528. [PubMed: 2391631]
18. Blazer, DG., George, LK. Established populations for epidemiologic studies of the elderly, 1996–1997: Piedmont health survey of the elderly, fourth in-person survey [Durham, Warren, Vance, Granville, and Franklin counties, North Carolina]. 2004. icpsr 2744 URL <http://www.icpsr.umich.edu/icpsrweb/NACDA/studies/2744>

Figure 1.

Empirical power of T_{CM} and T_F in Simulation A (no covariates) with $M = 25, 100$ and $\lambda = 2, 10$. Each entry is based on 500 Monte Carlo replicates.

Figure 2.
Patients' mean attachment loss (mm) plotted versus the number of teeth.

Figure 3. Empirical distribution functions of the attachment loss. The solid line is the conventional estimate, and the dashed line is the inverse-cluster-size weighted estimate.

Table 1

Empirical rejection probabilities (sizes) in the null case. The nominal significance level is set at 0.05. Each entry is based on 5000 Monte Carlo replicates.

Test						
Simulation	β	M	λ	T_F	T_{CM}	Kruskal-Wallis
A	0	25	2	0.041	0.075	0.089
		25	10	0.057	0.077	0.107
		50	2	0.034	0.065	0.074
		50	10	0.052	0.089	0.099
		100	2	0.031	0.061	0.069
		100	10	0.050	0.094	0.096
A	1	25	2	0.037	0.071	0.081
		50	2	0.032	0.064	0.070
		100	2	0.024	0.058	0.070
Simulation	β	M	a	T_F	T_{CM}	Kruskal-Wallis
B	0	25	$\Phi^{-1}(\frac{1}{3})$	0.043	0.054	0.058
		25	0	0.040	0.046	0.052
		25	$\Phi^{-1}(\frac{2}{3})$	0.044	0.049	0.053
B	0	50	$\Phi^{-1}(\frac{1}{3})$	0.037	0.042	0.047
		50	0	0.045	0.053	0.053
		50	$\Phi^{-1}(\frac{2}{3})$	0.042	0.047	0.051
B	0	100	$\Phi^{-1}(\frac{1}{3})$	0.038	0.045	0.049
		100	0	0.038	0.047	0.047
		100	$\Phi^{-1}(\frac{2}{3})$	0.044	0.049	0.050
B	1	25	0	0.044	0.050	0.051
		50	0	0.045	0.052	0.053
		100	0	0.037	0.044	0.046

Author Manuscript

Author Manuscript

Author Manuscript

Author Manuscript

Model A: Cluster sizes from Poisson-distribution; Y from a random intercept normal model.

Model B: Cluster sizes either 5 or 10 at proportions controlled by parameter α . Y with random intercept normal model and with errors from a $t(10)$ -distribution ($\nu_1 = \nu_2 = 10$).

$\Phi(\cdot)$ is the cumulative distribution function of the standard normal distribution.

Author Manuscript

Author Manuscript

Author Manuscript

Author Manuscript

Table 2
Empirical powers at Simulation B. Each entry is based on 1000 Monte Carlo replicates.

Configuration				Test			
<i>M</i>	ν_1	ν_2	γ	T_F	T_{CM}	<i>F</i> -test	Kruskal-Wallis
25	∞	1	0	0.071	0.099	0.000	0.038
	∞	3	0	0.049	0.055	0.045	0.054
	∞	10	0	0.048	0.052	0.058	0.057
	∞	∞	0.5	0.076	0.087	0.097	0.094
	∞	∞	1	0.154	0.189	0.201	0.193
100	∞	∞	2	0.494	0.546	0.572	0.554
	∞	1	0	0.271	0.796	0.000	0.047
	∞	3	0	0.065	0.083	0.048	0.059
	∞	10	0	0.038	0.042	0.042	0.045
	∞	∞	0.5	0.180	0.209	0.211	0.211
	∞	∞	1	0.553	0.626	0.640	0.641
	∞	∞	2	0.982	0.991	0.994	0.991

Table 3

Ordinary and weighted least squares regression on attachment loss with socioeconomic index and smoking as explanatory variables.

		Standard		
		Estimate	error	Z
Ordinary least squares	Intercept	2.68	0.08	32.2
	socioeconomic index	−0.54	0.19	−2.8
	smoking	−0.40	0.11	−3.8
Weighted least squares	Intercept	2.86	0.11	25.8
	socioeconomic index	−0.76	0.27	−2.8
	smoking	−0.34	0.14	−2.5