COMPUTER METHODS AND PROGRAMS IN BIOMEDICINE 83 (2006) 57–62

journal homepage: www.intl.elsevierhealth.com/journals/cmpb

Creating non-parametric bootstrap samples using Poisson frequencies

James A. Hanley[a,*], *Brenda MacGibbon*[b]

[a] *Department of Epidemiology, Biostatistics and Occupational Health, McGill University, Montreal, Que. H3A 1A2, Canada*
[b] *GERAD and Département de mathématiques, Université du Québec à Montréal, Canada*

ARTICLE INFO

Article history:
Received 15 December 2004
Received in revised form 9 April 2006
Accepted 11 April 2006

Keywords:
Bootstrap
Multinomial
Poisson
Weights

ABSTRACT

This article describes how, in the high-level software packages used by non-statisticians, approximate non-parametric bootstrap samples can be created and analyzed without physically creating new data sets, or resorting to complex programming. The comparable performance of this shortcut method, which uses Poisson rather than multinomial frequencies for the numbers of copies of each observation, is demonstrated theoretically by evaluating the bootstrap variance in an example where the classic estimator of the sampling variance of the statistic of interest has a known closed form. For sample sizes of 50 or more, bootstrap standard errors obtained by this shortcut method exceeded those obtained by the standard version by less than 1%. The proposed method is also evaluated in two worked examples, involving statistics whose sampling distribution is more complex. The second of these is also used to illustrate when one can and cannot use non-parametric bootstrap samples.

1. Introduction

The bootstrap method [1,2], is now widely used by statisticians. It employs re-sampling, rather than closed-form expressions derived from statistical theory, to estimate the sampling distribution of a statistical estimator, and to derive *p*-values and confidence intervals. The principle is straightforward, and very appealing. In the non-parametric version, the actual data set, containing n observations, is used to create a number (B) of perturbed data sets ('bootstrap samples') of size n, and the estimator in question is applied to each one; the variation among these B estimators is then used to describe its sampling variation. In the non-parametric version, the bootstrap sample consists of copies of the actual observations. In the parametric version, the bootstrap samples are generated from the parametrically-specified distribution that is assumed to have given rise to the original sample; the parameter values used to specify this distribution are estimated from the actual observations.

In practice, even the simpler non-parametric bootstrap technique is out of the reach of most non-statisticians—unless they have ready access to a statistician who is comfortable using the general bootstrap tools available in statistical packages such as S-Plus and R. Non-statisticians tend to use higher-level statistical packages such as Stata, SAS, and SPSS. These packages only have bootstrap inference tools available for a limited number of procedures. Of the three, Stata appears to have the broadest built-in implementation. There is just one specific implementation ('MULTTEST') in SAS, although with some programming ingenuity, it can be used for other applications. For both SAS and SPSS, end-users tend to rely on user-supplied macros, each one for a specific statistic, available from web pages of vendors, or user groups, or individuals located via an internet search. Of course, the more enterpris-

---

[*] *Corresponding author*. Tel.: +1 514 398 6270; fax: +1 514 398 4503.
E-mail address: James.Hanley@McGill.CA (J.A. Hanley).

0169-2607/$ – see front matter © 2006 Elsevier Ireland Ltd. All rights reserved.
doi:10.1016/j.cmpb.2006.04.006

ing end-users can modify a program, or macro, written for another estimator, to suit the one in question.

Interestingly, some authors, e.g. [3], have developed separate 'front-ends' to generate non-parametric bootstrap samples, supply each sample in turn to a specific 'stand-alone' statistical procedure, collect the parameter estimates, and summarize their distribution to calculate a bootstrap standard error or confidence interval. The aim was to "obviate the need for expensive, high-end statistical packages and associated script files written using in-built programming languages". Their method can be adapted to other statistical procedures used in biomedical modeling applications. Others, e.g. [4] have focused on a specific high-end statistical package and taken advantage of the fact that it has bootstrapping built into one specific routine. Hutson [5], although dealing with the specific problem of correlated (repeated measures) data, uses a data set of size $n = 4$ to explain how to create bootstrap samples for any statistical procedure. The key to his approach is to determine the (multinomial) number of copies of each original observation to place in a particular bootstrap sample. He shows how a random number function which generates realizations from different binomial distributions can be used to achieve this.

2. Background

For the past decade or so, one of the authors [JH] has programmed the non-parametric bootstrap calculations himself. Although the examples here will use SAS, the technique we describe is more general, and simpler. When the data set is small, the simplest approach is to create one larger data set containing $B \times n$ observations—with an additional variable (say 'sample') to identify each bootstrap sample. Each bootstrap sample consists of n observations, drawn with replacement, from the n actual observations. Then, using the 'BY sample' statement, one applies the SAS procedure (PROC) in question, and saves the relevant parameter estimates from each of the B bootstrap samples to a file. From this file, with B estimates, PROC UNIVARIATE can then be used to obtain the relevant features of the sampling distribution. When the data set is large, and storage space become an issue, it is more efficient to use a macro to create a bootstrap data set, run the PROC on it, append the relevant parameter estimates to a file, and delete this bootstrap sample before repeating the cycle. Our impressions of the relative speeds of the two approaches have been confirmed by the formal evaluation of Novikov [6] who found that the 'BY' approach was 80 times faster than the macro approach. On occasion, a mix of the two approaches is helpful: a macro can process k 'batches' of data sets; for each batch, a data set containing size B/k bootstrap samples is created, the 'BY' approach is used to obtain k estimates, and then this data set is deleted.

More recently, JH has been using a different system, provided the PROC in question allows one to use the FREQ statement. It can be explained as follows: let us consider an actual data set containing n observations and p variables, so that the data file is of size $n \times p$. Say that a particular bootstrap sample from this contains two 'copies' of actual observation 1, no 'copy' of actual observation 2, one of actual observation 3, etc. This sample can be analyzed directly using the actual data set, but with these 'numbers of copies' $(2, 0, 1, \ldots)$ used as frequencies in the FREQ statement (or 'weights' in Stata, or in SPSS). Thus, rather than create B new data sets of size $n \times p$, one can append the B sets/columns of frequencies—as B new variables, say $\text{freq}_1 \ldots \text{freq}_B$—to the actual data set, thereby creating a new data set of size $n \times (p + B)$. A simple macro can then be used to run the PROC in question B times, using each 'freq' variable in turn as the FREQ variable.

For a particular bootstrap sample, the challenge is how to generate a random frequency ('number of copies') to append to the ith observation in the actual data set. Technically, the set of n frequencies follows a multinomial distribution in which there are n independent 'trials', each with n outcome categories, and category probabilities $(1/n, \ldots, 1/n)$ for each trial [1, p. 286]. Although they can be generated exactly with some programming, there is a much simpler shortcut. Even though it does not yield frequencies that are exactly what is stipulated, we have been sufficiently reassured by preliminary tests to adopt it for routine use when sample sizes have been large.

In this note, the method is explained, and in a simple situation amenable to exact calculation, its performance is compared with that of the exact ('standard') version. An actual data set example which Efron and Tibshirani [1] used to illustrate the bootstrap algorithm for more complicated situations is also reworked, and used to compare the standard and proposed methods.

If a built-in bootstrap routine only reports standard errors and confidence limits, without showing the distribution of the B bootstrap estimates, end-users may not realize that the non-parametric bootstrap does not always work well, and that the parametric version may be more accurate. We are unable to offer omnibus methods for the much more sophisticated statistical programming needed for the parametric version. However, we do present a second worked example, where for one quantity of interest the non-parametric version works well, but for another, the parametric version is more appropriate.

3. Computational method: performance in a known situation

Instead of generating n random frequencies from a multinomial distribution with n equiprobable categories, and a sample size of n, one can generate n independent random frequencies from a Poisson distribution with expectation 1. Whereas in the standard version, the bootstrap samples are all of size n, this shortcut results in bootstrap samples of variable size: their expected size is n, but the S.D. is $n^{1/2}$.

To see how this affects the performance, consider one of the simplest of all inference calculations, the estimation of the sampling variability of the mean, $\bar{y}$, of n independently selected observations, $\{y_1, \ldots, y_n\}$, from a distribution with mean μ and variance σ^2. Denote the multinomial and Poisson frequencies by $\{m_1, \ldots, m_n\}$ and $\{p_1, \ldots p_n\}$, respectively.

Then, the standard bootstrap estimate is:

$$\bar{y}^* = \frac{\sum y_i \times m_i}{n} = \frac{\sum y_i \times m_i}{\sum m_i},$$

whereas the shortcut version, based on the Poisson frequencies, is:

$$\bar{y}^{**} = \frac{\sum y_i \times p_i}{n'} = \frac{\sum y_i \times p_i}{\sum p_i}.$$

Among (i.e. over) all possible bootstrap samples, the variance of the standard $\bar{y}^*$ is:

$$\mathrm{Var}[\bar{y}^*] = \frac{1}{n^2}\sum\sum y_i y_j \,\mathrm{Covar}[m_i, m_j]. \tag{1}$$

Now, $E[m_i] = n \times (1/n) = 1$, $\mathrm{Var}[m_i] = n \times (1/n) \times (1 - 1/n) = (n-1)/n$, and $\mathrm{Covar}[m_i, m_j] = n \times (0 - (1/n) \times (1/n)) = -1/n$. Substituting these into (1) yields:

$$\mathrm{Var}[\bar{y}^*] = \frac{1}{n^2}\sum y_i^2 (n-1)/n - \sum_i \sum_{j \neq i} y_i y_j (-1/n).$$

To simplify matters, without any loss of generality, take $\sum y_i = 0$, so that $\bar{y} = 0$; $\sum (y_i - \bar{y})^2 = \sum y_i^2$ and $\sum_i \sum_{j \neq i} y_i y_j = -\sum y_i^2$. Furthermore, in the usual notation, let $s^2 = \sum (y_i - \bar{y})^2/(n-1)$. Then, the variance in question simplifies to:

$$\begin{aligned}\mathrm{Var}[\bar{y}^*] &= \frac{1}{n^2}\{(n-1)/n + 1/n\}\sum y_i^2 = \frac{1}{n^2}\{1\}(n-1)s^2 \\ &= \frac{n-1}{n}\frac{s^2}{n} = \frac{n-1}{n}\widehat{\mathrm{Var}}[\bar{y}],\end{aligned} \tag{2}$$

where $\widehat{\mathrm{Var}}[\bar{y}]$ is the classic analytic estimator of the sampling variation. Davison and Hinkley [2, p. 22] establish this property by a more direct argument.

Since $n' = \sum p_i$ is variable, the overall variation, $\mathrm{Var}[\bar{y}^{**}]$, of the proposed Poisson-based estimator $\bar{y}^{**}$ can be calculated as:

$$\mathrm{Var}[\bar{y}^{**}] = \mathrm{Var}_{n'}[E[\bar{y}^{**}|n'\text{fixed}]] + E_{n'}[\mathrm{Var}[\bar{y}^{**}|n'\text{fixed}]], \tag{3}$$

where the inner E and Var are calculated conditionally—over all bootstrap samples with n' fixed—and the outer Var and E are evaluated using the distribution of n'. With n' fixed:

$$\begin{aligned}E_{n'}[\bar{y}^{**}|n'\text{fixed}] &= (1/n')\sum y_i \times E[p_i] = (1/n')\sum y_i \times (n'/n) \\ &= (1/n)\sum y_i = \bar{y},\end{aligned}$$

no matter the size of n', and so the first (variance) term in (3) is zero. With n' fixed, the approach that led to Eq. (2) allows us to write the conditional variance as:

$$V[\bar{y}^{**}|n'\text{fixed}] = \frac{1}{n'}\frac{n-1}{n}s^2,$$

and so the overall variance of $\bar{y}^{**}$ is the expectation of this quantity over the distribution of n'. Now, by construction, n' has a Poisson distribution with expectation n. Technically, it could take on a value of zero, but with the sample sizes to which this is applied ($n > 20$), we ignore this very unlikely scenario, and compute the expected variance numerically using $\mathrm{Prob}[n' = 1]$ onwards. Fig. 1 shows that the ratio $\mathrm{Var}[\bar{y}^{**}]/\mathrm{Var}[\bar{y}^*] = (1/n) \times E[1/n']$ is always greater than unity, but that the average over-estimation, relative to the standard bootstrap, is of the order of $(100/n)$%, for example a 2% overestimation of the variance when $n = 50$, a 1% overestimation when $n = 100$ and so on. Indeed, the ratios suggest that:

$$\mathrm{Var}[\bar{y}^{**}] \approx \{1 + 1/n\}\mathrm{Var}[\bar{y}^*] = \frac{n+1}{n}\frac{n-1}{n}\frac{s^2}{n} = \frac{n^2-1}{n^2}\widehat{\mathrm{Var}}[\bar{y}] \tag{4}$$

4. Worked examples

4.1. *Example 1. Principal components and a 'scientific intelligence quotient'*

Efron and Tibshirani [1, Chap. 7] used two examples to illustrate that the simple bootstrap algorithm "can provide standard errors for very complicated situations". In the first of these, they analyze the test scores of $n = 88$ students on five scientific subjects. These data are available at http://www.epi.mcgill.ca/hanley/software. The focus is on the structure of these multivariate data, and on the extent to which the five scores can be collapsed into just one, representing what the authors term a scientific "intelligence quotient". The extent to which the five scores represent a single dimension is usually measured by the proportion of variance explained by this first principal component. It is calculated as the ratio of the largest eigenvalue to the total of the eigenvalues, and in this example is $679.2/(679.2 + 199.8 + 102.6 + 83.7 + 31.8) = 0.619$. Efron and Tibshirani repeated the calculation on 200 bootstrap replications drawn from the original 88 subjects, to arrive at a bootstrap standard error of 0.047. Moreover, the histogram of the 200 bootstrap proportions "looked reasonably normal", suggesting that the 0.047 can be used to form a symmetric 95% CI of $0.618 \pm 1.96 \times 0.047$.

The SAS code for the standard bootstrap based on multinomial frequencies of inclusion and for our proposed use of Poisson frequencies to generate 'approximate' bootstrap

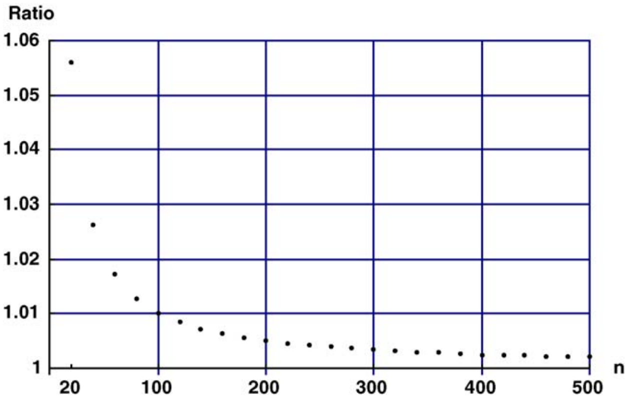

Fig. 1 – Ratio of the variance of the proposed bootstrap estimator to that of the standard bootstrap estimator of the sampling variability of the mean of *n* independent observations, as a function of *n*. The proposed estimator uses samples in which the number of times an observation is included follows a Poisson distribution with mean 1. The standard bootstrap uses frequencies, which follow a multinomial distribution with *n* equiprobable outcome categories, and *n* trials.

Table 1 – Estimates of standard error obtained by the standard (multinomial-based) and the proposed (Poisson-based) versions of the non-parametric bootstrap, applied to two data sets

Data set/parameter	Method	5000 samples	25 batches of 200	
			Median	Inter-quartile range
Efron/PC1[a]	Standard	0.0477	0.0489	0.0455–0.0496
	Proposed	0.0480	0.0480	0.0462–0.0502
Miller/ρ[b]	Standard	1.67	1.69	1.61–1.73
	Proposed	1.70	1.68	1.63–1.75

[a] PC1: first principal component (point estimate 0.619). Data are the test scores of 88 subjects, analyzed by Efron and Tibshirani [1, Table 7.1].

[b] ρ is the ratio of mortality rate before, to that after, change-point (point-estimate: 7.82). Data are the survival times of 184 patients post cardiac transplant [Miller and Halpern, 7].

samples are available online. In this case, the small size of the data set meant that it was just as practical to create a single data set containing all 200 bootstrap data sets as it was to use a single data set of 88, where each observation had 200 associated frequencies (the former has the advantage that it does not require a macro).

The user will find that the second significant digit of the standard error varies with the starting seed used for the random number generation. JH (somewhat lazily) uses the same seven digit odd number, e.g. 5555555 for many of his simulations, and it just so happened that using this initial seed to create the 200 samples gave a standard error of 0.0475, very close to that reported by Efron and Tibshirani. A seed of 3333333 gave an estimate of 0.0505. This illustrates that—with both the standard and proposed methods—the estimate of the standard error is itself stochastic, and that a larger number of bootstrap samples is required to produce estimated standard errors that are stable to two decimal places.

The top half of Table 1 compares the standard errors obtained by the two methods. Since the sample size is $n = 88$, the proposed method, as might have been expected, yields the same stability as the standard one.

4.2. *Example 2. Change-point: how long is the high-risk period after cardiac transplant?*

A goal in many survival analyses is to provide clinicians with brief summaries, which emphasize the important conclusions and are easy to communicate. Knowing the shapes of mortality and other event-rate functions, for example, can be an important for the management of patients, the scheduling of follow-up examinations, and when communicating with patients and relatives. Although the mortality rate in the days, months and years following heart transplant probably changes more gradually and more smoothly over this time, it may be more useful to present the mortality rate as a simpler function. A number of authors have used a simple change-point model for this reason. In this model, the mortality rate (or 'hazard' function) is constant (λ_1) up until a certain day τ, then changes to a lower constant (λ_2) beyond τ. This 'change-point' parameter τ, and the ratio $\rho = \lambda_1/\lambda_2$, which measures the size of the change in rates, have been the object of a number of data-analyses. Here we consider the estimation of these parameters from a classic data set, containing the post-transplantation survival of the 184 patients who received heart transplants in the Stanford heart transplantation program before February 1980 [7,8]. At the time of analysis, 113 of the patients had died, leaving 71 survival times right-censored.

The maximum likelihood estimators of these parameters were developed by Yao [9], but only for the case where none of the observations are censored. Approximate confidence intervals using the likelihood ratio process were constructed by Loader [10]. For confidence intervals for the change-point τ, Pham and Nguyen [11] showed that the parametric bootstrap could be used, but that the non-parametric one could not. In contrast, for confidence intervals for the rate ratio ρ, Yao's work suggested that both the parametric and non-parametric bootstrap methods would be appropriate—but only if samples sizes were 'large enough'. MacGibbon and Groshen [12] extended Yao's work and Pham and Nguyen's results concerning the parametric bootstrap to data sets that include censored observations.

Since the focus of this paper is on the non-parametric bootstrap, we use this change-point estimation problem primarily to evaluate how well our shortcut Poisson-based method performs in relation to the standard multinomial-based one for the rate ratio ρ. And as a by-product, we ask how the official non-parametric one compares to the parametric one. However we also use the example to show how the non-parametric bootstrap (both the standard and our proposed shortcut versions) can be easily misused to provide misleading results in the case of the more challenging estimation of the change-point parameter, τ.

The results are shown in Fig. 2. For the ratio parameter, the smooth distributions of parametric and non-parametric bootstrap estimates are virtually interchangeable, as are the standard and proposed versions of the latter (right half of Fig. 2, and bottom half of Table 1). However, for the change-point parameter, the problems with the non-parametric version (even the standard one) are quite evident in the multimodal and over-dispersed distributions (left half of Fig. 2). By contrast, the distribution of the parametric bootstrap estimates is quite regular, and agrees well with the range estimated by the likelihood-based one calculated by Loader [10].

5. Discussion

The theoretical calculations show that for sample sizes beyond 20, the 'Poisson frequencies' shortcut provides very similar variance estimates to the harder-to-program standard non-parametric bootstrap estimator. It also performs well in ac-

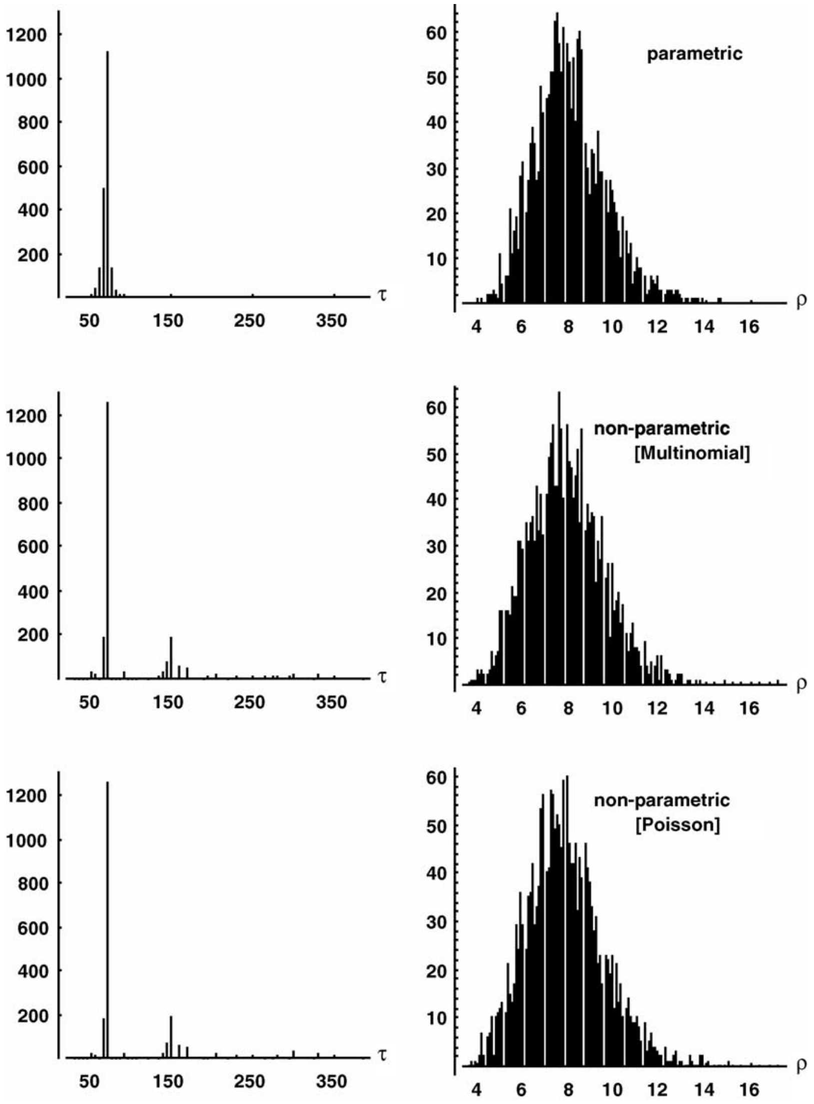

Fig. 2 – Frequency distributions of $B = 2000$ estimates of change-point (τ, days after transplant, shown in left column) and death rate ratio (ρ, right column) obtained by parametric and non-parametric bootstrap methods {the latter obtained by using multinomial frequencies—the standard way—and by Poisson frequencies—the shortcut way}.

tual data sets where the estimator has a complex form. Thus, we expect it also provides similar variance estimates for even more complex estimators.

Although Efron and Tibshirani—in problem 7.4 [1, p. 84]—ask the reader to calculate the Poisson approximation for the multinomial probability that any particular observation occurs exactly *k* times in a bootstrap sample, they do not exploit this approximation. We used this approximation in cell occupancy problems [13,14] but it also applies to the present 'numbers of copies' situation. While Davison and Hinkley never mention the Poisson distribution, their Table 2.2 [2, p. 23] does present the multinomial frequencies (in a simple example where $n = 10$) in a way that suggests the Poisson approximation for larger sample sizes.

Despite this, our search of the literature did not find any description or evaluation of the 'Poisson' shortcut described here. However, using keywords Poisson and bootstrap, we did find a description of a sequential bootstrap based on the Poisson distribution [15]. The authors' main objective was to create bootstrap samples, of variable size, in which there are always $m = n(1 - e^{-1}) = 0.632 \times n$ distinct observations. To do this, they use what they call a "Poisson bootstrap" approach, in which a set of independent Poisson frequencies, such as we have proposed, is used only if exactly *m* of these frequencies are non-zero; otherwise the sample of frequencies is rejected, and a new set of *n* independent Poisson frequencies is drawn. In previous work [16] they had shown that this sequential approach to resampling gave results which were within a distance $O(n^{-3/4})$ from those of the usual bootstrap.

Although not quite as rigorous, the method proposed here is much simpler to implement, and does not require knowledge of the sample size, *n*, in order to generate the frequencies. This same approach has been used in online bootstrap aggregating ('bagging') in which a random number (0, 1, 2, ...) of copies of each arriving observation needs to be generated as the observations appear, i.e. one by one, rather than in batch once all of the observations have arrived [17,18]. In all of our applications, the entire set of n observations is already available, but they enter the SAS data set one by one, and so it is easier to generate B independent Poisson frequencies for each one, without having to consider the frequencies for the other observations.

We did find one other work, [19] where this simpler approach was used. These authors note that for a fixed bootstrap sample of size n, the distribution of the inclusion frequencies (what we call $\{m_1 \ldots m_n\}$) is multinomial, but state that "it is simpler (and does not make a big difference—when n is sufficiently large)" if one works with a Poisson distribution for the size of the bootstrap sample, with n as the "mean number of data points in the sample". It was not clear whether this statement applied only to their theoretical calculations, or to the use of bootstrapping in practice.

The main purpose of this note is to convey a practical approach for end-users, particularly those who do their work in SAS or SPSS. In writing it, we discovered mainly theoretical research, spread out in some disconnected disciplines. We also came upon some very intuitive approaches, such as the Bayesian bootstrap [20] that sadly have not yet been exploited in statistical software for end-users.

Lastly, we do caution that there are some situations where the non-parametric bootstrap does badly. Inferences using the bootstrap are sensitive to model assumptions as illustrated by our change-point example and by Rubin's Bayesian bootstrap estimates of a proportion. The authors endorse his statement that "there are no general data analytic panaceas that allow us to pull ourselves up by our bootstraps". Although each problem must be treated separately, we can offer the following suggestions. Giné and Zinn [21] showed that the bootstrap method was usually valid when and only when the central limit theorem holds, so the bootstrap should perform well for large sample sizes when we know our estimates are asymptotically normal. However, in all cases histograms of the frequency distributions should be plotted and studied. Any evidence of multimodality or 'granularity' should indicate that the bootstrap is probably inappropriate in this situation. Our change-point example illustrates this point.

6. Template

The proposed method can be tailored to each user's statistical practice; thus, we do not provide an omnibus program. End-users may wish to choose and/or adapt one of the different approaches shown in the SAS code, depending on the size of their data set, and the capabilities of the statistical package they use.

Acknowledgments

This work was supported by grants from The Natural Sciences and Engineering Research Council of Canada and Le Fonds Québécois de la recherche sur la nature et les technologies.

Supplementary Material: SAS code

The supplementary material, in the form of SAS code, can be found, in the online version, at http://www.epi.mcgill.ca/hanley/software.

REFERENCES

[1] B. Efron, R.J. Tibshirani, An Introduction to the Bootstrap, Chapman & Hall, New York, 1993.
[2] A.C. Davison, D.V. Hinkley, Bootstrap Methods and Their Applications, Cambridge University Press, 1997.
[3] J. Parke, N.G.H. Holford, B.G. Charles, A procedure for generating bootstrap samples for the validation of non-linear mixed-effects population models, Comput. Methods Prog. Biomed. 59 (1999) 19–29.
[4] S.R. Cole, Simple bootstrap statistical inference using the SAS system, Comput. Methods Prog. Biomed. 60 (1999) 79–82.
[5] A.D. Hutson, A semiparametric bootstrap approach to correlated data analysis problems, Comput. Methods Prog. Biomed. 73 (2004) 129–134.
[6] I. Novikov, A remark on efficient simulations in SAS, J. Roy. Stat. Soc. D: The Statistician 52 (1) (2003) 83–86.
[7] R. Miller, J. Halpern, Regression with censored data, Biometrika 69 (1982) 521–531.
[8] D.R. Cox, D.O. Oakes, Analysis of survival data, Chapman & Hall, 1984.
[9] Y.-C. Yao, Maximum likelihood estimation in hazard rate models with a change-point, Commun. Stat. A: Theory and Methods 15 (1986) 2455–2466.
[10] C.R. Loader, Inference for a hazard rate change-point, Biometrika 78 (1991) 749–757.
[11] D.T. Pham, H.T. Nguyen, Bootstrapping the change-point of a hazard rate, Ann. Inst. Stat. Math. 45 (1993) 331–340.
[12] B. MacGibbon, S. Groshen, Confidence intervals for change-points in the exponential hazard rate model with censoring, in: Proceedings of the International Conference on Recent Advances in Survey Sampling, a Celebration of the Work of J.N.K. Rao, 2003, pp. 147–160.
[13] J.A. Hanley, Jumping to coincidences: defying odds "in the realm of the preposterous", Am. Stat. 46 (1992) 197–202.
[14] J.A. Hanley, B. MacGibbon, The Poisson distribution and "multiple events" cell occupancy problems, Biometrics 51 (1995) 1557–1561.
[15] G.J. Babu, P.K. Pathak, C.R. Rao, Second-order correctness of the Poisson bootstrap, Ann. Stat. 27 (5) (1999) 1666–1683.
[16] C.R. Rao, P.K. Pathak, V.I. Koltchinskii, Bootstrap by sequential resampling, J. Stat. Plann. Infer. 64 (1997) 257–281.
[17] N.C. Oza, S. Russell, Online bagging and boosting, in: T. Richardsons, T. Jaakkola (Eds.), Artificial Intelligence and Statistics, 2001, pp. 105–112.
[18] H.K.H. Lee, M.A. Clyde, Lossless online Bayesian bagging, J. Mach. Learn. Res. 5 (2004) 143–151.
[19] D. Malzahn, M. Opper, Approximate analytical bootstrap averages for support vector classifiers, in: S. Thrun, L. Saul, B. Schölkopf (Eds.), Advances in Neural Information Processing Systems, vol. 16, MIT Press, Cambridge, MA, 2004.
[20] D. Rubin, The Bayesian bootstrap, Ann. Stat. 9 (1) (1981) 130–134.
[21] E. Giné, J. Zinn, Bootstrapping general empirical measures, Ann. Probab. 18 (1990) 851–869.