This item was submitted to [Loughborough's Research Repository](#) by the author.
Items in Figshare are protected by copyright, with all rights reserved, unless otherwise indicated.

Subsampling bias and the best-discrepancy systematic cross validation

PLEASE CITE THE PUBLISHED VERSION

<https://doi.org/10.1007/s11425-000-0000-0>

PUBLISHER

Springer Verlag

VERSION

AM (Accepted Manuscript)

PUBLISHER STATEMENT

This is a post-peer-review, pre-copyedit version of an article published in Science in China Series A: Mathematics. The final authenticated version is available online at: <https://doi.org/10.1007/s11425-018-9561-0>

LICENCE

CC BY-NC-ND 4.0

REPOSITORY RECORD

Guo, Liang, Jianya Liu, and Ruodan Lu. 2019. "Subsampling Bias and the Best-discrepancy Systematic Cross Validation". Loughborough University. <https://hdl.handle.net/2134/38228>.

Subsampling bias and the best discrepancy systematic cross validation

Journal:	<i>SCIENCE CHINA Mathematics</i>
Manuscript ID	SCM-2018-0786.R1
Manuscript Type:	Original Article
Date Submitted by the Author:	10-Apr-2019
Complete List of Authors:	Guo, Liang; Shandong University at Weihai, Statistics Liu, Jianya; Shandong University, School of Mathematical Sciences Lu, Ruodan; University of Cambridge, Engineering
Keywords:	Subsampling Bias, Cross Validation, Systematic Sampling, Low Discrepancy Sequence, Best Discrepancy Sequence

SCHOLARONE™
Manuscripts

Subsampling Bias and The Best-Discrepancy Systematic Cross Validation

Liang Guo¹, Jianya Liu^{2,*} & Ruodan Lu^{3,4}

¹*Data Science Institute, Shandong University, Weihai 264209, China;*

²*Data Science Institute, Shandong University, Jinan 250100, China;*

³*School of Architecture, Building and Civil Engineering, Loughborough University, Loughborough LE11 3TU, UK;*

⁴*Darwin College, University of Cambridge, Cambridge CB3 9EU, UK*

Email:liangguo@sdu.edu.cn,jyliu@sdu.edu.cn, rl508@cam.ac.uk

Received ; accepted

Abstract Statistical machine learning models should be evaluated and validated before putting to work. Conventional k -fold Monte Carlo Cross-Validation (MCCV) procedure uses a pseudo-random sequence to partition instances into k subsets, which usually causes subsampling bias, inflates generalization errors and jeopardizes the reliability and effectiveness of cross-validation. Based on ordered systematic sampling theory in statistics and low-discrepancy sequence theory in number theory, we propose a new k -fold cross-validation procedure by replacing a pseudo-random sequence with a best-discrepancy sequence, which ensures low subsampling bias and leads to more precise *Expected-Prediction-Error (EPE)* estimates. Experiments with 156 benchmark datasets and three classifiers (logistic regression, decision tree and naïve bayes) show that in general, our cross-validation procedure can extrude subsampling bias in the MCCV by lowering the *EPE* around 7.18% and the variances around 26.73%. In comparison, the stratified MCCV can reduce the *EPE* and variances of the MCCV around 1.58% and 11.85% respectively. The Leave-One-Out (LOO) can lower the *EPE* around 2.50% but its variances are much higher than the any other CV procedure. The computational time of our cross-validation procedure is just 8.64% of the MCCV, 8.67% of the stratified MCCV and 16.72% of the LOO. Experiments also show that our approach is more beneficial for datasets characterized by relatively small size and large aspect ratio. This makes our approach particularly pertinent when solving bioscience classification problems. Our proposed systematic subsampling technique could be generalized to other machine learning algorithms that involve random subsampling mechanism.

Keywords Subsampling Bias, Cross Validation, Systematic Sampling, Low-Discrepancy Sequence, Best-Discrepancy Sequence

MSC(2010) 62-07; 11J71; 62G09; 68T05

Citation: Guo L, Liu J, Lu R. Sci China Math, 2019, 62, doi: 10.1007/s11425-000-0000-0

1 Introduction

All statistical machine learning processes start with a basic truism: garbage in, garbage out. This means that nonsense input data usually produce nonsense conclusions. However, “garbage in, garbage out” is not only about data quality, but also about the way in which data scientists evaluate and choose statistical machine learning models. Such evaluation often relies on a Cross-Validation (CV) to estimate model generalization performance metrics. The holdout procedure is the simplest kind of CV, in which a data set is split into a training set and a test set. However, the holdout is not reliable as the evaluation depends

*Corresponding author

heavily on which data points end up in the training set and which in the test set [24]. Thus, the evaluation may be significantly different depending on how the subsampling is made [24]. The k -fold Monte-Carlo CV (MCCV), the most popular CV procedure, is an effective way to reduce the subsampling bias in the holdout procedure. Given a dataset of n instances, an MCCV procedure is to randomly generate k subsets, each of which contains m instances, where $n = km$. That is to say, each instance is first assigned an index between 1 and n to construct a subsampling frame. Using the Simple Random Subsampling (SRS) method, the MCCV procedure then uses a Pseudo-Random Number Generator (PRNG) to generate a random integer sequence. This sequence consists of n distinct, randomly ordered integers between 1 and n , each of which refers to an element of the subsampling frame. Next, the random integer sequence is split into k equal-sized parts and the corresponding instances are partitioned into k subsets. Subsequently, k iterations of training and testing are performed such that in each iteration, a different subset is held out to estimate the *EPE* while the remaining $k - 1$ subsets are used for training [42].

However, relying on the SRS technique to subsample instances to training and test sets not only generates high computational costs, but also poses risks [5, 32]. The inductive bias of a model and the subsampling bias in the SRS may be confounded and lead to inflated *EPEs* [32]. This is especially true for small, imbalanced, and high dimensional datasets [16]. The random integer sequence used in the SRS poses a threat to the MCCV's validity and reliability, as the sequence may seem to be random, but not truly random and not truly uniformly distributed. This is because the elements in such sequence often clump at certain arbitrary intervals [17]. When performing SRS, similar instances could be included in one subset, meaning that information may be over- (clusters) or under-represented (voids) in that subset. This is just like the random shuffle function of *Spotify* (a music player App) cannot ensure songs from a given artist are spread throughout the playlist. To avoid the risk of obtaining dissimilar subsets, the MCCV procedure should be repeated p times (usually $p = 50$) with different random seeds to circumvent the risk of obtaining greatly dissimilar subsets [32], as there are trillions of ways to partition n instances into k subsets.

The reliability of the MCCV increases as the k and p increase. When $k = n$, the MCCV becomes the Leave-One-Out CV (LOO). In practice, the LOO is not widely used because it is more computationally demanding and has a larger variance than the MCCV [24]. The stratified MCCV procedure is another mean to alleviate the subsampling bias issue [40]. The stratified MCCV procedure starts by dividing a dataset into non-overlapping strata based on its class label vector. It then randomly and proportionally partitions instances into k subsets from each stratum. While this purposive sampling method ensures the same distribution of class values in all subsets, it may not ensure the same distributions of feature values. Thus, subsampling bias problem may remain unsolved. A low subsampling bias is crucial because no matter how sophisticated a statistical machine learning algorithm is, the resulting model may be biased if its training sets fail to represent the idiosyncrasies of the complete set [5, 32]. Likewise, subsampling bias and insufficiently informative test data will inflate the *EPEs*, jeopardizing the validity and reliability of a CV procedure [16]. We may run the risk of choosing a biased model and drawing a wrong conclusion.

This study follows the Ordered Systematic Subsampling (OSS) theory [30] to propose a simple yet effective method, the best-discrepancy systematic subsampling method for k -fold CV (hereafter the BDSCV), which ensures a smaller subsampling bias than the MCCV, stratified CV and LOO. The contribution of our proposed systematic subsampling method lies in the fact that we make a step further to reduce subsampling bias than the conventional systematic subsampling one. That is, the first subsampling instance is not chosen at random as in the conventional OSS method. Instead, it is chosen based on the Best-Discrepancy Sequence (BDS), a sequence that is a uniformly distributed sequence, finite or infinite, in which the number of any proportion of its elements falling into an arbitrary interval proportionally converts to the distance of the interval [33]. Our subsampling interval is also determined by the BDS.

This article is organised as follows. We first introduce the theories of the low discrepancy sequence (LDS) and the BDS. We then illustrate subsampling traps in conventional k -fold CV. Next, we propose the best-discrepancy systematic subsampling method and the BDSCV procedure. We conduct experiments to provide empirical supports to our theories. Finally, we discuss our proposed procedure and its limitations.

2 Methods

2.1 Low-discrepancy sequences and best-discrepancy sequences

For brevity, we restrict our discussion to the key definitions of low-discrepancy sequence (LDS) and best-discrepancy sequence (BDS), which are necessary to understand our procedure. We refer interested readers to the comprehensive reviews of [27], [21] and [14].

Let $\mathcal{X}$ denote a sequence $x_1, x_2, \dots, x_n, \dots$ in the unit interval $[0, 1)$. The discrepancy D_N of its consecutive N elements is defined by

$$D_N = \sup_{0 \leq a < b < 1} \left| \frac{C([a, b); N)}{N} - (b - a) \right|, \tag{2.1}$$

where the counting function $C([a, b); N)$ denotes the number of element x_n with $1 \leq n \leq N$ for which $x_n \in [a, b)$ [27]. This $\mathcal{X}$ is said to be an LDS if D_N is low for all sufficiently large integers N . The sequence $\mathcal{X}$ is uniformly distributed in $[0, 1)$ if $D_N \rightarrow 0$ as $N \rightarrow \infty$ [27]. In simple terms, the sequence $\mathcal{X}$ is uniformly distributed if every subinterval $[a, b)$ eventually obtains its proper share in $[0, 1)$. A sequence's level of uniformity is measured by its discrepancy function D_N . A classical theorem of Chung [9] states that, in general, the PRNG does not produce a sequence with D_N lower than $N^{-1/2}$. In any case, the order of the discrepancy D_N cannot be lower than $N^{-1} \log N$ by [35]. Thus, the best we can hope for D_N is something like

$$D_N \leq cN^{-1+\varepsilon} \tag{2.2}$$

where $\varepsilon > 0$ is arbitrarily small, and $c > 0$ is some constant. This explains the following definition.

Definition 2.1. A sequence $\mathcal{X}$ is called a BDS, if its discrepancy function D_N satisfies (2.2) for arbitrary $\varepsilon > 0$ and some constant $c > 0$.

In the following we will present some BDSs that are produced by methods from number theory. Of particular interests to this paper are those BDSs of the form

$$\{n\alpha\}, \quad n = 1, 2, 3, \dots \tag{2.3}$$

where $\{\alpha\}$ denotes the fractional part of a real number α . Plainly each term of (2.3) falls into the unit interval $[0, 1)$. Denote by $\mathcal{X}(\alpha)$ the sequence (2.3).

Theorem 2.2. The sequence $\mathcal{X}(e)$ is a BDS, where e is the base of natural logarithm.

Proof. It follows from [1] that (2.2) holds for $\mathcal{X}(e)$ with the constant c therein depending only on ε , and by [31] the dependence of c on ε is effectively computable. See also the arguments in [21, Chap. 4, §5]. □

Remark 2.3. Of course the sequence $\mathcal{X}(\alpha)$ can be a BDS for many other α . For example, it is known [21, Chap. 4, §5] that if α is real algebraic, then the sequence $\mathcal{X}(\alpha)$ is a BDS. Examples of these α include $\sqrt{p}$ with p ranging over all prime numbers p . Prior studies also propose other ways to generate BDS, for example the Halton sequence [20] and the van der Corput sequence [10]. Note that our sequence $\mathcal{X}(e)$ in Theorem 2.2 belongs to none of the above examples since e is transcendental. In this paper we have followed the suggestion of Hua and Wang [21, Chap. 4, §5] to use $\mathcal{X}(e)$ rather than those previous sequences, and our numerical experiments support their claim that $\mathcal{X}(e)$ is of important practice implications.

The discrepancy properties of LDS/BDS have attracted a lot of attentions [15]. Indeed, point sampling is an important research area in computational mathematics. For example, Xu and Zhou [46] propose a deterministic method to produce the interpolation points and show its outstanding performance with $\mathcal{L}_1$ minimization. Previous studies also report that the LDS can accelerate the convergence for Monte Carlo quadrature formulas [33] and its points are correlated to provide greater uniformity [36]. The LDS has been applied in many fields other than mathematics, including optimization [18, 26, 34, 44], machine

learning [2, 3, 8], computer graphics and vision [23, 25, 28, 37, 45], finance [4, 39, 41, 43], and engineering [7, 12, 22, 29]. Our study is inspired from these studies, although they do not directly address CV procedures. We develop the best-discrepancy systematic subsampling method and the BDSCV procedure based on the BDS.

2.2 Numerical remarks on Section 2.1

Panel A-C in Figure 1 shows the scatter plots of three random sequences, each of which contains 500 elements. The three sequences contain clear clumps and voids so that it is highly possible that the distance between two consecutive random elements may be either very small or very large. To formally compute the D_N of each sequence, we randomly choose 50 pairs of interval $[a, b)$ and compute the discrepancy to uniform distribution for these two sequences using (2.1) respectively. Results are shown in Figure 2, from which we find that the average D_N and the variance of 50 measures of each random sequence are quite large. That is to say, these three sequences are not quite uniformly distributed.

Figure 1 Here.

Figure 2 Here.

Panel D in Figure 1 illustrates a BDS of 500 elements. The interval between two consecutive elements of a BDS is determined by the property of uniformity so that the elements scatter evenly to avoid clumps and voids. We formally compute the D_N of BDS by randomly choosing 50 pairs of interval $[a, b)$ using (2.1). Figure 2 shows that the resulting discrepancy of BDS is close to zero and the variance of the 50 measures of BDS is greatly smaller than that of its random counterpart. That is, unlike the elements of a random number sequence, the elements of a BDS are as uniformly distributed as possible. Successive elements are inserted in a position as far as possible from the other elements in order to avoid clustering and voids. The elements generated sequentially fill in the larger gaps between the previous elements of the sequence [8].

Random sequences cause subsampling bias in k -fold CV. To illustrate subsampling traps, we generate an artificial dataset of 16 instances with four binary features and one label vector with 16 classes (i.e. each combination of the values of the four binary features corresponds to one class). We duplicate this dataset five times and form a complete set of 80 instances. A logistic regression model is used to map out the relations between features and labels. This should be an easy task, as the classification boundary of this dataset are perfectly clear. Therefore, if the 80 instances are well partitioned into representative subsets, then the error rate of both Holdout and CV procedure should be zero.

We conduct k -fold CV with the logistic regression model, where k ranges from two to 10. Each k -fold CV procedure is repeated 50 times with different random seeds to circumvent the risk of obtaining largely dissimilar subsets. Results are shown in Figure 3. The error rate of the k -fold CV is surprisingly high and the error rate drops when the number of k increase. However, the error rates are always larger than zero. The box plots also suggest that there are large variations over 50 repetitions. This illustration vividly demonstrates that there are actually traps in conventional k -fold CV. A major threat to the validity and reliability of k -fold CV lies in subsampling bias by the clumps and voids illustrated in Figure 1, which largely jeopardize the representativeness of subsets. Results of model evaluation and selection could be problematic.

Figure 3 Here.

2.3 The Best-Discrepancy Systematic Subsampling Method

We propose a systematic subsampling method based on the uniformity property of the BDS. The theoretical foundation of our proposed method stems from sampling theory, which proves that the ordered systematic subsampling (OSS) method can obtain smaller sampling error than the SRS and the stratified sampling [30]. The OSS method works as follows. A complete dataset should be firstly sorted either ascendingly or descendingly so that adjacent elements tend to be more similar than elements that are

farther apart. Then, a fixed subsampling interval (also called a “skip”) k is defined as $k = n/m$, where n is the dataset size and m is the subset size. The subsampling starts by selecting an element from the ordered complete dataset at random and then every k^{th} element is selected until a subset of m elements is formed. The ordered systematic subsampling outperforms the SRS because the former can avoid choosing two elements adjacent to each other and the resulting subsets are forced to compose of various values, meaning that the discrepancy between subsets caused by subsampling bias can be reduced [30].

We improve the OSS by replacing the fixed subsampling interval with a BDS. Suppose $D = (d_1, \dots, d_n)$ is a one-dimensional dataset with n instances. We use this vector notation to emphasise the importance of the order of D . Let $\tilde{D} = (\tilde{d}_1, \dots, \tilde{d}_n)$ be the sorted D , meaning that $\tilde{D}$ is a permutation of the coordinates of D , so that $\tilde{d}_1 < \dots < \tilde{d}_n$. This $\tilde{D}$ is to be divided into k subsets, each of which contains m instance (i.e. $n = km$). To do so, an integer sequence derived from a BDS is needed that refers to the subsampling frame of a dataset.

Let $X = (x_1, \dots, x_n)$ be the vector of the first n elements of $\mathcal{X}(e)$. Let $\tilde{X} = (\tilde{x}_1, \dots, \tilde{x}_n)$ be the sorted X , meaning that $\tilde{X}$ is a permutation of the coordinates of X , so that $\tilde{x}_1 < \dots < \tilde{x}_n$. The ranking vector $R = (r_1, \dots, r_n)$ of X is such that $\tilde{x}_{r_j} = x_j$ for $j = 1, \dots, n$. Then, given $n = km$, we can chunk R into k disjoint subsets of the same size m , that is

$$R = \bigcup_{1 \leq j \leq k} R_j \tag{2.4}$$

where $R_1 = (r_1, \dots, r_m)$, $R_2 = (r_{m+1}, \dots, r_{2m})$, etc. This R inherits the property of uniformity. We can use R to partition $\tilde{D}$ into k subsets $\tilde{D} = \cup_{1 \leq j \leq k} \tilde{D}_j$, where for $j = 1, \dots, k$,

$$\tilde{D}_j = (\tilde{d}_r : r \in R_j). \tag{2.5}$$

Note that each subset $\tilde{D}_j$ is written as a vector for the same reason.

Our proposed best-discrepancy systematic subsampling method can make sure the instances in a subset as heterogeneous as possible while different subsets as homogenous as possible. According to the theory of analysis of variance [30], we know that the *Sum of Squares of Total (SSTO)* is decomposed into the *Sum of Squares Within subsets (SSW)* plus the *Sum of Squares Between subsets (SSB)*. The level of heterogeneity of instances within subsets can be measured by the *Intra-class Correlation Coefficient (ICC)* [30],

$$ICC = 1 - \frac{km}{km - 1} \frac{SSW}{SSW + SSB}, \tag{2.6}$$

where

$$SSW = \sum_{i=1}^k \sum_{j=1}^m (d_{ij} - \bar{d}_i), \tag{2.7}$$

$$SSB = n \sum_{i=1}^k (\bar{d}_i - \bar{d}), \tag{2.8}$$

$$SSTO = SSW + SSB. \tag{2.9}$$

And Equation (2.6) can be transformed as

$$SSW = \frac{(1 - ICC)(n - 1)SSTO}{n}. \tag{2.10}$$

Equation (2.10) implies that SSW becomes larger when ICC becomes negative (i.e. the instances within a subset are dispersed more than a randomly chosen subset would be). Prior study [30] proves that to have a negative ICC , the dataset should be ordered, so that it exhibits a positive autocorrelation

(i.e. subsequent instances have similar or the same values) and instances should be chosen with a skip from the ordered dataset so that a subset is forced to compose of both small values and large values. The best-discrepancy systematic subsampling method follows exactly these guidelines to allocate instances. In addition, it is important to mention that the subsampling interval is naturally and uniformly set by the uniformity property of a BDS. The subsampling interval of the BDS guarantees that a subset is composed of heterogeneous instances.

2.4 Numerical remarks on Section 2.3

We illustrate the relations among the BDS, ICC and SSW in the following example. Suppose a one-dimensional dataset D is given, and the 21 instances of D are to be partitioned into three subsets with (2.5). We use $\mathcal{X}(e)$ to generate a tiny BDS using e 's 38 digits after the decimal place (see the first column of Table 1). The property of uniformity of the BDS ensures that the distribution of values in a subset should be the same or close to that of the complete set, and the between-subset variance should equal or be close to zero. For example, the second column of Table 1 demonstrates that the variances of the three subsets ($\sigma_1^2 = \sigma_2^2 = \sigma_3^2 = 0.08$) are all close to that of the complete set $\sigma_X^2 = 0.08$, provided that the complete set is evenly chunked into three subsets using the BDS. The third and fourth columns of Table 1 show that the ranking vector R of the BDS exhibits the same uniformity property as the means of the three subsets are quite close (10, 10, and 12 respectively).

Then, D is ascendingly sorted as $\tilde{D}$ so that adjacent instances tend to be more similar than instances that are farther apart. The 21 instances are partitioned into three subsets with (2.5) (see the fifth to seventh columns of Table 1). The three subset variances are quite close to that of D ($\sigma_D^2 = 0.01$) and the SSB is almost zero (0.0082). As expected, the ICC is negative (-0.1242) and the SSW is 0.2168. The instances with adjunct or the same values (for example, the 9th, 10th and 11th elements of D) are forced to spread out between different subsets. The ranking vector R makes the distributions of values highly similar among three subsets, as none of them consists of unrepresentative (i.e. mainly low or high) values.

To compare, we use the SRS to split D into three subsets and repeat the procedure 50 times. The averages of SSB of these SRS procedures is 0.0208, which is 2.531 times larger than that of the best-discrepancy systematic subsampling method. This experiment shows that subsampling bias can be greatly reduced by the proposed method, even though the sample size of this tiny dataset is only 21. We conduct an additional experiment with 500 datasets of different sizes (ranging from 54 to 999) and with different levels of variance (ranging from 0 to 0.025), skewness (ranging from 0 to 0.116), and kurtosis (ranging from 0.506 to 31.398). The results suggest that regardless the characteristics of a dataset, the SSB of the best-discrepancy systematic subsampling method is much smaller (23.16%) than that of the SRS.

Table 1 Here.

In short, this example proves that the best-discrepancy systematic subsampling method can satisfy Lohr's [30] two conditions that make the ICC negative. In contrast, the SRS cannot guarantee all subsets are of heterogeneous instances. Indeed, there are too many possibilities to randomly generate k subsets. It is then difficult to avoid "bad", quasi-homogeneous subsets with only or mainly small or large values. Therefore, the best-discrepancy systematic subsampling procedure outperforms the SRS in terms of subset representativeness.

3 The BDSCV Procedure

We propose a k -fold CV procedure that is built upon the best-discrepancy systematic subsampling procedure. The pseudocode of BDSCV is shown in Figure 4.

Figure 4 Here.

Suppose we are given an independently and identically distributed m dimensional dataset D of n instances, where $D \in \mathbb{R}^{n \times m}$. Each instance in D consists of a class label vector Y and vectors of measured feature covariates X . As discussed in the previous section, D is sorted before subsampling. However, it is difficult to decide along which axis to sort D . An arbitrary yet effective way is to transform D from multi-dimensional space to one-dimensional space using a dimensionality-reduction technique H , which produces a linear transformation $T \in \mathbb{R}^{m \times 1}$ and results in $E = DT$, where $E \in \mathbb{R}^{n \times 1}$ and E is the underlying structure of D and points to a direction where there is the most variance and where the elements in E are most spread out.

Simply stated, it is just like to project all dimensions of a dataset on the component vector E , which is a combination of all dimensions with weights. The instances with close or the same values in E are considered similar in the terms of all dimensions but not just of label classes. And E can be served as an axis along which D is sorted so that instances that are similar to some extent are then gathered and separated from dissimilar ones. Therefore, we ascendingly sort E as $\tilde{E}$ and then sort D as $\tilde{D}$, so that $\tilde{E} = \tilde{D}T$. These $\tilde{E}$ and T are then discarded and will not be included in subsequent analyses.

Next, we generate a BDS with n elements and then a ranking vector R accordingly. We partition $\tilde{D}$ into k subsets, using the same allocation mechanism in (2.5).

Thirdly, the k subsets of $\tilde{D}$ are used to conduct a standard k -fold CV procedure with a classifier A . Let F_k be the k -th fold that serves as the test set, and $\tilde{D}_k$ be the training set obtained by removing the instances in F_k from $\tilde{D}$. The estimated EPE over k folds with the dimensionality-reduction technique H is estimated as in [32]:

$$EPE(\tilde{D})_{H,A} = \frac{1}{K} \sum_{k=1}^K \widehat{EPE}(F_k). \tag{3.1}$$

And the variance of the CV estimator is estimated as in [32]:

$$\sigma_{H,A}^2 = \frac{1}{K} \sum_{k=1}^K (EPE(F_k)_{H,A} - EPE(\tilde{D})_{H,A})^2. \tag{3.2}$$

Finally, different dimensionality-reduction techniques may lose different amounts of information from the complete dataset [11]. Sweeping different candidate techniques, such as Principal Component Analysis (PCA), Factor Analysis (FA), kernel PCA (kPCA), or truncated Singular Value Decomposition (tSVD), will enable us to fine tune the BDSCV procedure, so that the original space of D can be reduced to one-dimensional space while retaining the most important information as much as possible in E . Therefore, previous steps are repeated using different dimensionality-reduction techniques in order to find the best H in terms of $EPE(\tilde{D})_{H,A}$ and its corresponding $\sigma_{H,A}^2$ as the final CV results for the classifier A .

4 Experiments

We first test the logistic regression model in Section 2.2 with the proposed 10-fold BDSCV. The model achieves zero error rate in every fold, which suggests that the BDSCV can eliminate subsampling bias to a great extent. We then test the BDSCV with the Penn Machine Learning Benchmark (PMLB) datasets [38]. Note that we eliminate 10 PMLB datasets that are too big to be computed on a personal computer. The characteristics of these 156 datasets are summarized in Table 2. We compare the performance of our proposed BDSCV to MCCV, stratified MCCV and LOO.

Table 2 Here.

The MCCV and stratified MCCV procedures are conducted using the *scikit-learn* cross-validation module (version 0.19.2) with 50 repetitions for the three classifiers respectively. Our proposed BDSCV is mainly built on the same *scikit-learn* module except the subsampling part. We use the MCCV as the benchmark and compute the following three comparative performance metrics as:

(a) the comparative *EPE* ratio of a classifier A , which is

$$\varphi EPE_A = \frac{EPE_{BDSCV/Stratified/LOO_A}}{EPE_{MCCV_A}}; \quad (4.1)$$

(b) the comparative variance ratio of a classifier A , which is

$$\varphi \sigma_A^2 = \frac{\sigma_{BDSCV/Stratified/LOO_A}^2}{\sigma_{MCCV_A}^2}; \quad (4.2)$$

and (c) the comparative computational burden ratio of the k -fold CV procedures, which is

$$\varphi Time = \frac{Time_{BDSCV/Stratified/LOO}}{Time_{MCCV}}. \quad (4.3)$$

Our experiments are conducted on a desktop personal computer (CPU: Intel Core i7-7700 eight cores, Memory: 64G, and SSD: 256GB). Note that computing time may slightly vary depending on the overall load and specifications of a computer. Table 3 shows that the number of datasets in which each of the four CV procedures achieve the same lowest *EPE*s and variances (it is possible that multiple CV procedures obtain the lowest performance metrics in a dataset). As expected, the BDSCV achieves the lowest *EPE*s and variances in most datasets, followed by the LOO and the stratified MCCV. Table 3 suggests that there are significant subsampling biases among different CV procedures and if we count on the most widely used 10-fold MCCV to do model selections, then we probably make the wrong choices.

Table 3 Here.

We summarize the three comparative performance metrics in Table 4. The *EPE*s, variances and computational time of each classifier and of each dataset can be found in Supplementary Material.

The comparisons between the BDSCV and the MCCV are as follows (see Panel A in Table 4). On average, the $EPE_{Logistic}$ and $\sigma_{Logistic}^2$ of the former is 94.04% and 70.30% of that of the latter respectively; the $EPE_{DecisionTree}$ and $\sigma_{DecisionTree}^2$ of the BDSCV is 88.02% and 77.50% of that of the MCCV respectively; and the $EPE_{Naïve Bayes}$ and $\sigma_{Naïve Bayes}^2$ of the former is 96.41% and 72.00% of that of the latter respectively. On average, the BDSCV can lower the *EPE*s around 7.18% and the variances around 26.73% regardless classifiers. In comparison, the stratified MCCV can slightly reduce subsampling bias by decreasing the *EPE*s of the MCCV around 1.58% and the variances around 11.85%. The LOO can lower the *EPE*s around 2.50% but its variances are much higher than the any other CV procedure (see Panel B and C in Table 4). The computational time of the BDSCV is just 8.64% of the MCCV, 8.67% of the stratified MCCV and 16.72% of the LOO (see Panel D in Table 4). These results suggest that the proposed BDSCV outperforms the other CV procedures as it can extrude significant subsampling bias while using a small fraction of computational time of the other CV procedures. Being computationally efficient is especially important where it is expensive or impractical to train multiple models.

Finally, our proposed CV procedure may not be a cure-all and be more pertinent to a dataset with particular characteristics. We compute the aspect ratio (i.e. the number of features divided by the number of instances) of the 156 datasets and divide them into four groups based on the quantiles of the aspect ratio. The first group includes 67 datasets whose number of features is smaller or equal to 1% of number of instances. The second group includes 20 datasets whose percentages of number of features to number of instances are in the range of (1%, 2%]. The third group includes 33 datasets whose aspect ratios are in the range of (2%, 5%]. And the fourth group include 36 datasets whose aspect ratios are larger than 5%. Table 4 shows that the BDSCV can significantly reduce subsampling bias in the third

and fourth groups of datasets, as the *EPEs* of the BDSCV are significantly smaller. These findings imply that the BDSCV could be particularly effective in analysing in bioscience datasets, which usually are characterized as high dimensionality and relatively small sample sizes [13, 16].

In short, these experiments confirm that the MCCV's suffers a lot from subsampling bias. The BDSCV can reduce subsampling bias in the *EPEs* and variances of the three classifiers to a significant extent. The improvements of the stratified MCCV are not as substantial as the BDSCV. The LOO does not have significant advantages compared to the other CV procedures. Both the stratified MCCV and the LOO are not as computationally efficient as the BDSCV. In addition, the decision tree is the most sensitive classifier to subsampling bias in terms of *EPEs*, especially in the datasets with high aspect ratios (i.e. the 3rd and 4th groups) in which the BDSCV can deflate the $EPE_{DecisionTree}$ around 15% on average. The logistic regression is the most sensitive classifier to subsampling bias in terms of variances, especially in the datasets with relatively low aspect ratio (i.e. the 1st, 2nd and 3rd groups) in which the BDSCV can reduce variances around 34%.

5 Discussion and Closing Remarks

A statistical machine learning model could be affected by the idiosyncrasies of training data if the discrepancy between *k* subsets is too high. As a result, the *EPEs* and variances of a MCCV procedure are contaminated by subsampling bias and cannot reflect the true generalization performance of a model's inductive mechanism. The inflated metrics will interfere with data scientists making right decisions. Our experimental results confirm that subsampling bias in the MCCV procedure greatly jeopardizes its validity and reliability for model selection even after 50 repetitions. Our experiments also confirm that the subsampling bias can be reduced through Lohr's guidelines [30]— systematic subsampling from an ordered population with a skip. The BDSCV follows these guidelines so that it is relatively more precise, stable and computationally efficient than all the other three CV procedures.

It is important to point out that the BDSCV is not meant to completely replace other CV procedures. As the results of our experiments suggest that the advantages of the different types of CV may depend on dataset characteristics and decision boundary. We recommend using the BDSCV in the datasets characterized with high aspect ratios (i.e. small dataset size and high dimensionality) and for the non-parametric classifier such as the decision tree. Besides, the BDSCV is much more computational efficient than the other CV procedures so it is suitable for the time-critical statistical machine learning problems. Finally, it is beneficial to estimate the true generalization performance of a learning function using different CV procedures in order to choose the model with the lowest *EPEs* and variances. Further studies are also necessary to determinate what is the most appropriate CV procedure for a particular dataset and for a statistical machine learning model.

In conclusion, this study sheds more light on the literature of CV and statistical machine learning model evaluation by proposing the BDSCV procedure, which is a new way to estimate a model's relatively objective generalizability by significantly reducing subsampling bias in existing CV procedures. It also helps the model to pick out the hidden underlying mapping between input-output vectors and to remain stable when the training set changes, leading to more precise estimations and more confident inferences. What is more, our proposed procedure is straightforward and computationally efficient. Unlike other information-preserving subsampling techniques [19] that need to understand the complete dataset in great detail, our procedure does not make the CV computation more complex. Lastly, our subsampling method may replace the random mechanism used in many statistical machine learning algorithms.

The Python source codes, Supplementary Material, and the 156 benchmark datasets are available at <http://cssci.wh.sdu.edu.cn/bdscv>

Acknowledgements The authors would like to thank the referees for helpful suggestions on an earlier draft of the paper. The second author is supported by the National Science Foundation of China under Grant 11531008, the Ministry of Education of China under Grant IRT16R43, and the Taishan Scholar Project of Shandong Province.

References

- 1 Baker A. On Some Diophantine Inequalities Involving the Exponential Function. *Canadian Journal of Mathematics*, 1965, 17: 616–626
- 2 Bergstra J, Bardenet R, Bengio Y, Kégl B. Algorithms for Hyper-parameter Optimization. *Advances in Neural Information Processing Systems*, 2011, 1: 2546–2554
- 3 Bergstra J, Bengio Y. Random Search for Hyper-parameter Optimization. *Journal of Machine Learning Research*, 2012, 13: 281–305
- 4 Boyle P, Broadie M, Glasserman P. Monte Carlo Methods for Security Pricing *Journal of Economic Dynamics and Control*, 1997, 21: 1267–1321
- 5 Braga-Neto U M, Dougherty E R. Is Cross-Validation Valid for Small-Sample Microarray Classification? *Bioinformatics*, 2004, 20: 374–380
- 6 Braga-Neto U M, Zolnari A, Dougherty G. Cross-Validation under Separate Sampling: Strong Bias and How to Correct It. *Bioinformatics*, 2014, 30: 3349–3355
- 7 Branicky M, LaValle S, Olson K, Yang L. Quasi-Randomized Path Planning. *Proceedings of IEEE International Conference on Robotics and Automation*, 2001, 2: 1481–1487
- 8 Cheng J. Computational Investigation of Low-Discrepancy Sequences in Simulation Algorithms for Bayesian Networks. *Proceedings of the Sixteenth Conference on Uncertainty in Artificial Intelligence*, 2000, 1: 72–81
- 9 Chung K. An Estimate Concerning the Kolmogoroff Limit Distribution. *Transactions of the American Mathematical Society*, 1949, 67: 36–50
- 10 van der Corput J G. Verteilungsfunktionen (Erste Mitteilung) *Proceedings of the Koninklijke Akademie van Wetenschappen te Amsterdam (in German)*, 1935, 38: 813–821
- 11 Cunningham J, Ghahramani Z. Linear Dimensionality Reduction: Survey, Insights, and Generalizations. *Journal of Machine Learning Research*, 2015, 16: 2859–2900
- 12 Dai H, Wang W. Application of Low-Discrepancy Sampling Method in Structural Reliability Analysis. *Structural Safety*, 2009, 31: 155–164
- 13 Díaz-Uriarte R, DeAndrés S. Gene Selection and Classification of Microarray Data Using Random Forest. *BMC Bioinformatics*, 2006, 7: 1–13
- 14 Dick J, Kuo F, Sloan I. High-Dimensional Integration: The Quasi-Monte Carlo Way. *Acta Numerica*, 2013, 22: 133–288
- 15 Dick J, Pillichshammer F. The weighted star discrepancy of Korobovs sets. *Proceedings of the American Mathematical Society*, 2015, 143: 5043–5057
- 16 Fu W, Carroll R, Wang S. Estimating Misclassification Error with Small Samples via Bootstrap Cross-Validation. *Bioinformatics*, 2005, 21: 1979–1986
- 17 Gentle J. *Statistics and Computing Random Number Generation and Monte Carlo Methods* New York: Springer, 2003
- 18 Georgieva A, Jordanov I. A Hybrid Meta-Heuristic for Global Optimisation Using Low-Discrepancy Sequences of Points. *Computers and Operations Research*, 2010, 37: 456–469
- 19 Groot P De, Postma G, Melssen W, Buydens L. Selecting a Representative Training Set for the Classification of Demolition Waste Using Remote NIR Sensing. *Analytica Chimica Acta*, 1999, 392: 67–75
- 20 Halton J H. Algorithm 247: Radical-inverse quasi-random point sequence *Communications of the ACM*, 1964, 7: 701–702
- 21 Hua L K and Wang Y. *Applications of Number Theory in Approximate Analysis*. Science Press, Beijing: 1978.
- 22 Kalagnanam J, Diwekar U. An Efficient Sampling Technique for Off-Line Quality Control. *Technometrics*, 1997, 39: 308–319
- 23 Keller A. The Fast Calculation of Form Factors Using Low Discrepancy Sequences. *Proceedings of the 12th Spring Conference on Computer Graphics*, 1996, 1: 195–204
- 24 Kohavi R. A Study of Cross-Validation and Bootstrap for Accuracy Estimation and Model Selection. *Proceedings of International Joint Conferences on Artificial Intelligence*, 1995, 14: 1137–43
- 25 Kollig T, Keller A. Efficient Multidimensional Sampling. *Computer Graphics Forum*, 2002, 21: 557–563
- 26 Kucherenko S, Sytsko Y. Application of Deterministic and Low-Discrepancy Sequence in Global Optimisation. *Computational Optimization and Applications*, 2005, 30: 297–318
- 27 Kuipers L, Niederreiter H. *Uniform Distribution of Sequences* New York: John Wiley & Sons, 1974
- 28 Li X, Wang W, Martin R, Bowyer A. Using Low-Discrepancy Sequences and the Crofton Formula to Compute Surface Areas of Geometric Models. *Computer Aided Design*, 2003, 35: 771–782
- 29 Lindermann R, Steven S, LaValle M. Incremental Low-Discrepancy Lattice Methods for Motion Planning. *Proceedings of IEEE International Conference on Robotics and Automation*, 2003, 1: 2920–2927
- 30 Lohr S. *Sampling: Design and Analysis* Boston: Brooks/Cole, 2009
- 31 Mahler, K., On a paper by A. Baker on the approximation of rational powers of e , *Acta Arith.* 27(1975), 61–87.
- 32 Molinaro A, Simon R, Pfeiffer R. Prediction Error Estimation: A Comparison of Resampling Methods. *Bioinformatics*, 2005, 21: 307–330

33 Niederreiter H. Random Number Generation and Quasi-Monte Carlo Methods. Philadelphia: SIAM, 1992

34 Pant M, Thangaraj R, Grosan C, Abraham A. Improved Particle Swarm Optimization with Low-Discrepancy Sequences. *Proceedings of the IEEE Congress on Evolutionary Computing*, 2008, 2: 3011–3018

35 Schmidt W. Irregularities of Distribution VII. *Acta Arithmetica*, 1972, 21: 45–50

36 Struckmeier J. Fast Generation of Low-Discrepancy Sequences. *Journal of Computational and Applied Mathematics*, 1995, 61: 29–41

37 Quinn J, Langbein F, Martin R. Low-Discrepancy Sampling of Meshes for Rendering. *Eurographics Symposium on Point-Based Graphics*, 2007, 1: 19–28

38 Olson R, LaCava W, Orzechowski P, Urbanowicz R, Moore J. PMLB: A Large Benchmark Suite for Machine Learning Evaluation and Comparison. *BioData Mining*, 2017, 10: 36

39 Paskov S, Traub J. Faster Valuation of Financial Derivatives. *The Journal of Portfolio Management*, 1995, 22: 113–123

40 Pedregosa F, Varoquaux G, Gramfort A, Vanderplas J. Scikit-Learn: Machine Learning in Python. *Journal of Machine Learning Research*, 2011, 12: 2825–2830

41 Singhee A, Rutenbar R. From Finance to Flip Flops: A Study of Fast Quasi-Monte Carlo Methods from Computational Finance Applied to Statistical Circuit Analysis. *Proceedings of the 8th International Symposium on Quality Electronic Design*, 2007, 1: 685–692

42 Stone M. Cross-Validatory Choice and Assessment of Statistical Predictions. *Journal of the Royal Statistical Society Series B*, 1974, 36: 111–147

43 Tan K, Boyle P. Applications of Randomized Low Discrepancy Sequences to the Valuation of Complex Securities. *Journal of Economic Dynamics & Control*, 2000, 24: 1747–1782

44 Uy N, Hoai N, McKay R, Tuan P. Initialising PSO with Randomised Low-Discrepancy Sequences: The Comparative Results. *Proceedings of the IEEE Congress on Evolutionary Computing*, 2007, 1: 1985–1992

45 Wenzel L, Dair D, Vazquez N. Pattern Matching System and Method with Improved Template Image Sampling Using Low Discrepancy Sequence. Washington: U.S. Patent No. 6,229,921, 2001

46 Xu Z Q, Zhou T. On sparse interpolation and the design of deterministic interpolation points. *SIAM Journal on Scientific Computing*, 2014, 36: A1752–A1769

12

Guo L *et al.* *Sci China Math for Review*

Figure 1 The Scatter Plots of Random Sequences and of BDS

Figure 2 The Uniform Discrepancy of Random Sequences and of BDS

Figure 3 Error Rate in Conventional k-fold CV

```
Read a dataset  $D$  of  $m$  dimensions and  $N$  instances
EPEs=[ ], VARs=[ ]
for  $H \leftarrow$  [PCA, FA, kPCA, tSVD] do:
    Use  $H$  to product a one-dimensional vector  $E$  from  $D$ 
     $E$  is attached to  $D$  as the last column
     $D$  is ascendingly sorted along  $E$  as  $\tilde{D}$ 
     $E$  is deleted from  $\tilde{D}$ 

    Generate a BDS  $X$  of  $N$  elements using Equation (3)
    Produce  $X$ 's ranking vector  $R$ 
     $R$  is chunked to  $k$  subsets
     $\tilde{D}$  is partitioned into  $k$  subsets using Equation (4)

    Conduct  $k$ -fold CV with a classifier  $A$  and get  $k$  EPEs
    Obtain the mean of  $k$  EPEs using Equation (11)
    Save the mean EPE in EPEs
    Obtain the variance among  $k$  EPEs using Equation (12)
    Save the variance in VARs
Choose the smallest mean EPE from EPEs
Choose the corresponding variance from VARs
```

Figure 4 Pseudocode of the BDSCV Procedure

Table 1 Elements with the Best Discrepancy Systematic Subsampling Method

BDS		R		$\tilde{D}$		
X_I	$subset_{X_I}$	R_I	$subset_{R_I}$	$\tilde{D}_I$	$\tilde{d}_i$ after partition	$subset_{d_i}$
0.7183		16		1.50 ₁	1.70 ₁₆	
0.4366		10		1.53 ₂	1.67 ₁₀	
0.1548	I	4	I	1.55 ₃	1.60 ₄	I
0.8731	$\mu_1: 0.44$	19	$\mu_1: 10$	1.60 ₄	1.80 ₁₉	$\mu_1: 1.66$
0.5914	$\sigma_1^2: 0.08$	13	$\sigma_1^2: 36$	1.62 ₅	1.68 ₁₃	$\sigma_1^2: 0.01$
0.3097		7		1.63 ₆	1.64 ₇	
0.028		1		1.64 ₇	1.50 ₁	
0.7463		17		1.65 ₈	1.76 ₁₇	
0.4645		11		1.66 ₉	1.67 ₁₁	
0.1828	II	5	II	1.67 ₁₀	1.62 ₅	II
0.9011	$\mu_2: 0.47$	20	$\mu_2: 10$	1.67 ₁₁	1.90 ₂₀	$\mu_2: 1.69$
0.6194	$\sigma_2^2: 0.08$	14	$\sigma_2^2: 36$	1.68 ₁₂	1.69 ₁₄	$\sigma_2^2: 0.01$
0.3377		8		1.68 ₁₃	1.65 ₈	
0.0559		2		1.69 ₁₄	1.53 ₂	
0.7742		18		1.70 ₁₅	1.78 ₁₈	
0.4925		12		1.70 ₁₆	1.68 ₁₂	
0.2108	III	6	III	1.76 ₁₇	1.63 ₆	III
0.9291	$\mu_3: 0.50$	21	$\mu_3: 12$	1.78 ₁₈	1.92 ₂₁	$\mu_3: 1.70$
0.6474	$\sigma_3^2: 0.08$	15	$\sigma_3^2: 36$	1.80 ₁₉	1.70 ₁₅	$\sigma_3^2: 0.01$
0.3656		9		1.90 ₂₀	1.66 ₉	
0.0839		3		1.92 ₂₁	1.55 ₃	
SSW	1.6678		756		0.2168	
SSB	0.0109		14		0.0082	
ICC	-0.1591		-0.1455		-0.1241	

Table 2 Descriptive Statistics of 156 PMLB Datasets

	#Instances	#Features	#Labels	Aspect Ratio
Minimum	32	2	2	0
Maximum	28056	240	26	0.54
Mean	2056.34	24.77	3.83	0.05
Median	2000	9	2	0.01

Table 3 Number of Datasets that Each CV Reaches the Lowest EPEs and Variances

Classifier	<i>BDSCV</i>		<i>MCCV</i>		<i>Stratified MCCV</i>		<i>LOO</i>	
	<i>EPE</i>	<i>Var</i>	<i>EPE</i>	<i>Var</i>	<i>EPE</i>	<i>Var</i>	<i>EPE</i>	<i>Var</i>
L^1	118	101	8	15	14	43	23	2
DT^2	115	102	8	21	12	43	43	12
NB^3	125	94	4	8	8	55	22	2

¹ L: Logistic
² DT: Decision Tree
³ NB: Native Bayes

Table 4 Comparisons between Different CV Procedures

A: BDSCV VS MCCV						
Aspect Ratio	EPE_L	σ_L^2	EPE_{DT}	σ_{DT}^2	EPE_{NB}	σ_{NB}^2
Q1	97.95%	64.29%	90.51%	87.88%	98.61%	75.11%
Q2	94.45%	67.06%	90.25%	64.22%	96.57%	57.31%
Q3	91.37%	66.80%	82.92%	78.72%	95.89%	72.47%
Q4	92.36%	83.06%	88.41%	79.17%	94.59%	83.11%
Average	94.04%	70.30%	88.02%	77.50%	96.41%	72.00%
B: Stratified MCCV VS MCCV						
Aspect Ratio	EPE_L	σ_L^2	EPE_{DT}	σ_{DT}^2	EPE_{NB}	σ_{NB}^2
Q1	99.48%	72.02%	98.91%	91.23%	99.59%	80.99%
Q2	97.96%	83.79%	99.53%	93.58%	99.91%	89.41%
Q3	97.25%	90.41%	94.35%	93.08%	98.79%	85.73%
Q4	98.40%	94.25%	97.62%	93.73%	99.25%	85.59%
Average	98.28%	85.12%	97.60%	92.90%	99.39%	86.43%
C: LOO VS MCCV						
Aspect Ratio	EPE_L	σ_L^2	EPE_{DT}	σ_{DT}^2	EPE_{NB}	σ_{NB}^2
Q1	99.72%	44036.96%	93.16%	36397.24%	100.01%	37667.87%
Q2	100.20%	7897.06%	92.52%	7693.12%	100.36%	7000.07%
Q3	97.43%	4977.84%	91.25%	4084.19%	101.24%	6016.48%
Q4	98.25%	4387.51%	98.45%	4411.73%	97.44%	3858.06%
Average	98.90%	15324.84%	93.85%	13146.57%	99.76%	13635.62%
D: Computational Time Comparisons						
Aspect Ratio	$\frac{BDSCV}{MCCV}$	$\frac{Stratified}{MCCV}$	$\frac{LOO}{MCCV}$	$\frac{BDSCV}{Stratified}$	$\frac{LOO}{Stratified}$	$\frac{BDSCV}{LOO}$
Q1	9.15%	99.62%	886.01%	9.19%	889.44%	3.08%
Q2	8.48%	99.68%	152.33%	8.51%	153.16%	9.68%
Q3	8.52%	99.79%	126.15%	8.54%	126.61%	20.60%
Q4	8.42%	99.80%	90.95%	8.43%	91.23%	33.53%
Average	8.64%	99.72%	313.86%	8.67%	315.11%	16.72%

L: Logistic

DT: Decision Tree

NB: Native Bayes