



LUDWIG-
MAXIMILIANS-
UNIVERSITÄT
MÜNCHEN

INSTITUT FÜR STATISTIK



Mathias Fuchs, Roman Hornung, Riccardo De Bin, Anne-Laure Boulesteix

A U -statistic estimator for the variance of resampling-based error estimators

Technical Report Number 148, 2013
Department of Statistics
University of Munich

<http://www.stat.uni-muenchen.de>



A U -STATISTIC ESTIMATOR FOR THE VARIANCE OF RESAMPLING-BASED ERROR ESTIMATORS

M. FUCHS, R. HORNUNG, R. DE BIN, A.-L. BOULESTEIX

ABSTRACT. We revisit resampling procedures for error estimation in binary classification in terms of U -statistics. In particular, we exploit the fact that the error rate estimator involving all learning-testing splits is a U -statistic. Therefore, several standard theorems on properties of U -statistics apply. In particular, it has minimal variance among all unbiased estimators and is asymptotically normally distributed. Moreover, there is an unbiased estimator for this minimal variance if the total sample size is at least the double learning set size plus two. In this case, we exhibit such an estimator which is another U -statistic. It enjoys, again, various optimality properties and yields an asymptotically exact hypothesis test of the equality of error rates when two learning algorithms are compared. Our statements apply to any deterministic learning algorithms under weak non-degeneracy assumptions. In an application to tuning parameter choice in lasso regression on a gene expression data set, the test does not reject the null hypothesis of equal rates between two different parameters.

Unbiased Estimator; Penalized Regression Model; U -Statistic; Cross-Validation; Machine Learning;

1. INTRODUCTION

The goal of supervised statistical learning is to develop prediction rules taking the values of predictor variables as input and returning a predicted value of the response variable. A prediction rule is typically learnt by applying a learning algorithm M to a so-called learning data set. A typical example in biomedical research is the prediction of patient outcome (e.g. recidive/no recidive within five years, tumor class, lymph node status, response to chemotherapy, etc.) based on bio-markers such as, e.g., gene expression data. The practitioners are usually interested in the accuracy of the prediction rule learnt from their data set to predict future patients, while methodological researchers rather want to know whether the learning algorithm is good at learning accurate prediction rules for different data sets drawn from a distribution of interest. The first perspective is called “conditional” (since referring to a specific data set) while the latter, which we take in this paper, is denoted as “unconditional”. Precisely, this paper focuses on the parameter defined as the difference between the unconditional errors of two learning algorithms of interest, M and M' , for binary classification.

If the data set is very large, one can observe independent realizations of estimators of the unconditional error rates and use them for a paired t -test (see Section 2.3). In practise, however, huge data sets are rarely available. Prediction errors are thus usually estimated by resampling procedures consisting of splitting the available data set into learning and test sets a large number of times and averaging the estimated error over these iterations. The well-known cross-validation procedure can be seen as a special case of resampling procedure for error estimation. A detailed overview of the vast literature on cross-validation would go beyond the scope of this paper. The reader is referred to Arlot & Celisse (2010) for a comprehensive survey.

Having estimated the error rate, it is typically of interest to test the null hypothesis of equal error rates between learning algorithms. This requires insight into the estimator's variance. Resampling-based error estimators typically have a very large variance, in particular in the case of small samples or high-dimensional predictor space (Dougherty et al., 2011). The estimation of this variance has been the focus of a large body of literature, especially in the machine learning context. A good estimation of the variance of resampling-based error estimators would allow to, e.g., derive reliable confidence intervals for the true error or to construct statistical tests to compare the performance of several learning algorithms. The latter task is of crucial importance in practise, since applied computational scientists including biostatisticians often have to make a choice within a multitude of different learning algorithms whose performance in the considered settings is poorly explored. In van der Wiel et al. (2009), for each splitting in repeated sub-sampling the predictions of the two classifiers are compared by a Wilcoxon-test, and the resulting p -values are combined. In Jiang et al. (2008a), the authors show the asymptotic normality of the error rate estimator in the case of a support vector machine. In Jiang et al. (2008b), a bias-corrected bootstrap-estimator for the error rate from leave-one-out cross-validation is introduced.

Various estimators of the variance of resampling-based error estimators have been suggested in the literature (Dietterich, 1998; Nadeau & Bengio, 2003), most of them based on critical simplifying assumptions. As far as cross-validation error estimates are concerned, Bengio & Grandvalet (2003) show that there exists no unbiased estimator of their variance. To date, the estimation of the variance of resampling-based error rates in general remains a challenging issue with no adequate answer yet both from a theoretical and practical point of view. In particular, there are no exact nor even asymptotically exact test procedures for testing equality of error rates between learning algorithms available. The present paper shows that there is an asymptotically exact test for the comparison of learning algorithms by using and extending results from U -statistics theory.

Despite the large body of literature, there seems to be no explicit treatment of the asymptotic properties of the estimators in general. Our main results are Theorem 3.9, stating that there is an unbiased estimator of the difference estimator's variance if $n \geq 2g + 2$, where g is the learning set size and n is the sample size, and Theorem 4.1, providing the central limit theorem for the studentized statistic as it is needed for testing. The use of only half the sample size for learning has already occurred in the literature, in a roughly similar context and on grounds of intuitive reasoning (Bühlmann & van de Geer, 2011, Section 10.2.1).

Corollary 5.1 gives an explicit bound on the number of iterations necessary to approximate the leave- p -out estimator, i.e. the minimum variance estimator, to an arbitrary given precision, where $p = n - g$; we show that this minimal variance can be estimated by an unbiased estimator, namely that from Definition 3.7. It has minimal variance itself, and the ensuing studentized test in (17) is asymptotically exact. This shows that it is not necessary to endeavour in determining the distribution of combinations of p -values to test equality of error rates, as in van der Wiel et al. (2009, Section 2.3).

Section 2 recalls important definitions pertaining to U -statistics and cross-validation viewed as incomplete U -statistics. We show that the procedure which involves all learning-testing splits is then a complete U -statistic. In Section 3, we show that U -statistics theory naturally suggests an unbiased estimator of the variance of this estimated difference of errors as soon as the sample size criterion is satisfied. In Section 4, we exploit this variance estimator to derive an asymptotically exact hypothesis test of equality of the true

errors of two learning algorithms M and M' . Section 5 addresses numerical computation of approximations of the leave- p -out cross-validation estimator, while an illustration of the variance estimation and the hypothesis test through application to the choice of the penalty parameter in lasso regression is presented in Section 6.

2. DEFINITIONS, NOTATIONS AND PRELIMINARIES

2.1. Hoeffding's definition. Since the complete error estimator that we will consider from the next section on is a U -statistic, we start by recalling the definition of U -statistics and their basic properties. In the following, the reader who is already familiar with the machine learning context may take $\Psi_0 := \Phi_0$ and $m := g + 1$ at first, where g is the learning sample size and Φ_0 is as defined in (5); however, this is not necessary and we will need other cases of the definition as well.

Definition 2.1 (U -statistic, Hoeffding (1948)). Let (\mathbf{Z}_i) , $i = 1, \dots, n$ be independent and identically distributed r -dimensional random vectors with arbitrary distribution. Let $m \leq n$, and let $\Psi_0 : \mathbb{R}^{r \times m} \rightarrow \mathbb{R}$ be an arbitrary measurable symmetric function of m vector arguments. Write $\Psi_0(S)$ as above for the well-defined value of Ψ_0 at those Z_i with indices from $S \subset \{1, \dots, n\}$, $|S| = m$. Consider the average of Ψ_0 in the maximal design \mathcal{S} of unordered size- m subsets

$$(1) \quad U = U(\mathbf{Z}_1, \dots, \mathbf{Z}_n) = \binom{n}{m}^{-1} \sum_{S \in \mathcal{S}} \Psi_0(S).$$

Any statistic of such a form is called a U -statistic.

The trailing factor is the inverse of the number of summands, the cardinality $|\mathcal{S}|$. So, a U -statistic is an unbiased estimator for the associated parameter

$$(2) \quad \Theta(P) = \int \cdots \int \Psi_0(\mathbf{z}_1, \dots, \mathbf{z}_m) dP(\mathbf{z}_1) \cdots dP(\mathbf{z}_m)$$

for a probability distribution P on \mathbb{R}^r , where $\mathbf{z}_1, \dots, \mathbf{z}_m$ are r -vectors. A parameter of this form is called a regular parameter. If m is the smallest number such that there exists such a symmetric function Ψ_0 that represents a given parameter $\Theta(P)$ in the form (2), then m is called the degree of Ψ_0 or of $\Theta(P)$, and the function Ψ_0 is called a kernel of U . Furthermore, U has minimal variance among all unbiased estimators over any family of distributions P containing all properly discontinuous distributions. Hoeffding (1948) shows the asymptotic normality of U -statistics as $n \rightarrow \infty$ and that a good number of well-known statistics such as sample mean, empirical moments, Gini coefficient, etc. are subsumed by the definition.

It is also possible to associate a U -statistic to a non-symmetric kernel Ψ , i.e. to estimate $\mathbb{E}(\Psi)$ by a U -statistic. The cost is to deal with $n!/(n-m)!$ summands, many more than just the binomial coefficient $n!/(m!(n-m)!)$. The symmetrization, indeed, consists in grouping all $m!$ summands involving the same unordered index set together. This writes a U -statistic with non-symmetric kernel Ψ in the form of Hoeffding (1948, Section 4.4) with symmetric kernel Ψ_0 as in Hoeffding (1948, Section 3.3).

2.2. The difference of true error rates as the parameter of interest. The goal of this section is to formalize the true error rate and to recall its nature as an expectation. Let $\mathcal{X} = \mathbb{R}^{r-1}$, $r \in \mathbb{N}$, be a fixed predictor space, and $\mathcal{Y} \subset \mathbb{R}$ be a space of responses. The number r will be as in Hoeffding (1948) and one would usually denote $p = r - 1$. Assume there is an unknown but fixed probability distribution P on $\mathcal{X} \times \mathcal{Y} \subset \mathbb{R}^r$ defined on the σ -algebra of Lebesgue-measurable sets. We do not require P to be absolutely continuous with respect

to the Lebesgue measure in order to allow for a discrete marginal distribution in \mathcal{Y} as it occurs in binary classification. The distribution P can be thought of as being supported on \mathbb{R}^r instead of $\mathcal{X} \times \mathcal{Y} \subset \mathbb{R}^r$ only, by identifying P with its push-forward image $i_*(P)$ under the inclusion $i: \mathcal{X} \times \mathcal{Y} \rightarrow \mathbb{R}^r$. This allows to apply Hoeffding (1948) which only describes U -statistics on a Euclidean space \mathbb{R}^r to the definition and investigation of U -statistics on (products of) $\mathcal{X} \times \mathcal{Y}$. Let us fix a loss function $L: \mathcal{Y} \times \mathcal{Y} \rightarrow \mathbb{R}$. Typically, L is the misclassification loss $L(y_1, y_2) = 1_{y_1 \neq y_2}$, but can be an arbitrary measurable function. Since we suppose the marginal distribution of P on \mathcal{Y} to be discrete, the loss function associated as done below to a learning algorithm is almost surely bounded. Therefore, all moments exist, which will be helpful throughout the paper. It is automatic for the misclassification loss. However, some of the following also work for an unbounded distribution on \mathcal{Y} . In that case one would typically consider losses which are not almost surely bounded, such as, for instance, the residual sum of squares loss $L(y_1, y_2) = (y_1 - y_2)^2$ or the Brier score in survival analysis. We will not work out the details of unbounded losses.

Say we are interested in the difference of error rates of classifiers learnt on a sample of size g . Typical choices are $g = 4n/5$ (assuming that five divides n) for a learning/testing sample size ratio of 4 : 1 and $g = n - 1$ for leave-one-out cross-validation. We also allow for $g = 0$ in case we are interested in the performance of classification rules that were already learnt on different and fixed data. In that case, it is important that the learning data were different since otherwise there would be a problematic contradiction between simultaneously regarding the data as fixed and as being drawn from P .

Let $(\mathbf{z}_i)_{i=1, \dots, n} = (\mathbf{x}_1, y_1, \dots, \mathbf{x}_n, y_n)$, where $\mathbf{x}_i \in \mathcal{X}$ and $y_i \in \mathcal{Y}$, and denote by

$$f_M: (\mathcal{X} \times \mathcal{Y})^{\times g} \times \mathcal{X} \rightarrow \mathcal{Y}$$

the function that maps $(\mathbf{z}_1, \dots, \mathbf{z}_g; \mathbf{x}_{g+1}) \in (\mathcal{X} \times \mathcal{Y})^{\times g} \times \mathcal{X}$ to the prediction, an element of \mathcal{Y} , by the learning algorithm M , learnt on $\mathbf{z}_1, \dots, \mathbf{z}_g$ and applied to \mathbf{x}_{g+1} . We are only concerned with deterministic learning algorithms M , i.e. ones which do not involve any random component for classification. We suppose that f_M is symmetric in the first g entries $\mathbf{z}_1, \dots, \mathbf{z}_g$, i.e. M treats all learning observations equally, and that f_M is measurable with respect to the product σ -algebra. The inclusion $\mathcal{X} \times \mathcal{Y} \rightarrow \mathbb{R}^r$ defines an inclusion $(\mathcal{X} \times \mathcal{Y})^{\times g} \times \mathcal{X} \rightarrow \mathbb{R}^{rg+r-1}$ and in order to be able to apply Hoeffding (1948) we view f_M as a map on \mathbb{R}^{rg+r-1} by extending it by zero on $\mathbb{R}^{rg+r-1} \setminus (\mathcal{X} \times \mathcal{Y})^{\times g} \times \mathcal{X}$ which is a null-set with respect to the push-forward measure $i_*(P)$. The map f_M is then also measurable on \mathbb{R}^{rg+r-1} .

Denote by $\Phi: (\mathcal{X} \times \mathcal{Y})^{g+1} \rightarrow \mathbb{R}$ the function

$$(3) \quad \Phi(\mathbf{z}_1, \dots, \mathbf{z}_g; \mathbf{z}_{g+1}) := L(f_M(\mathbf{z}_1, \dots, \mathbf{z}_g; \mathbf{x}_{g+1}), y_{g+1}) - L(f_{M'}(\mathbf{z}_1, \dots, \mathbf{z}_g; \mathbf{x}_{g+1}), y_{g+1})$$

for two learning algorithms M and M' . The value $\Phi(\mathbf{z}_1, \dots, \mathbf{z}_g; \mathbf{z}_{g+1})$ is the empirical difference of error rates between M and M' , learnt on the first g observations $(\mathbf{z}_1, \dots, \mathbf{z}_g)$ of the sample (\mathbf{z}_i) and evaluated on the single last entry. The semicolon thus visually separates learning and test sets. The definition of Φ involves only a single test observation. Anticipating a little, the reason is that Φ had to be defined with a minimal number of arguments necessary for (6) below. In the following, we will conveniently consider larger test sample sizes by applying the mechanism of associating a U -statistic to a kernel.

As noted above, we assume that Φ is almost surely bounded. This happens, for instance, for bounded loss functions such as the misclassification loss. Also, we may view Φ as being defined on $\mathbb{R}^{r(g+1)}$ instead of $(\mathcal{X} \times \mathcal{Y})^{g+1}$ without notational distinction, in the same way as f_M was extended to \mathbb{R}^{rg+r-1} .

The true difference of error rates between M and M' is the expectation of Φ , taken with respect to $g + 1$ independent realizations of P :

$$\begin{aligned}
 \Delta &:= e_M - e_{M'} = \mathbb{E}_{P^{\otimes(g+1)}}(\Phi(\mathbf{Z}_1, \dots, \mathbf{Z}_g; \mathbf{Z}_{g+1})) \\
 (4) \quad &= \int \cdots \int_{(\mathcal{X} \times \mathcal{Y})^{\times(g+1)}} (L(f_M(\mathbf{z}_1, \dots, \mathbf{z}_g; \mathbf{x}_{g+1}), y_{g+1}) \\
 &\quad - L(f_{M'}(\mathbf{z}_1, \dots, \mathbf{z}_g; \mathbf{x}_{g+1}), y_{g+1})) dP(\mathbf{z}_1) \cdots dP(\mathbf{z}_{g+1}),
 \end{aligned}$$

where both learning and test data are random. The existence of the expectation follows from measurability and from the boundedness assumption. The quantity Δ is the parameter of main interest. Also, we consider the symmetric function of $g + 1$ arguments

$$(5) \quad \Phi_0(\mathbf{z}_1, \dots, \mathbf{z}_{g+1}) := \frac{1}{g+1} \sum_{i=1}^{g+1} \Phi(\mathbf{z}_1, \dots, \mathbf{z}_{i-1}, \mathbf{z}_{i+1}, \dots, \mathbf{z}_{g+1}; \mathbf{z}_i),$$

satisfying

$$(6) \quad \Delta = \mathbb{E}(\Phi) = \mathbb{E}(\Phi_0).$$

For the particular non-symmetric kernel $\Psi = \Phi$ introduced in (3), the symmetrization $\Psi_0 = \Phi_0$ of (5) can be written involving only $m = g + 1$ summands instead of $m!$, in other words only cyclic permutations instead of all permutations, due to the assumption that learning is symmetric. In practise, it is not advantageous to compute Φ_0 directly because a learning procedure should be used on more than just one test observation for numerical efficiency (see Section 5); however, it is very convenient to consider Φ_0 for ease of presentation.

Remark 2.2. In case one is interested in estimating e_M only instead of a difference $\Delta = e_M - e_{M'}$, one can set the second summand of (3) identically to zero. We will not go into the details.

2.3. Tests of the true error rate. In this section, let us recall the test problem of interest. We want to test the null hypothesis $H^0 : \mathbb{E}(\Phi) = 0$ against the alternative $H^1 : \mathbb{E}(\Phi) \neq 0$. The former is usually called the unconditional null hypothesis (Braga-Neto & Dougherty, 2004).

Remark 2.3. There is also a conditional null hypothesis where the classification rule is supposed to be given, for instance learnt on fixed independent data, and the expectation is taken only with respect to the test set. However, the learning data are usually also random and may even be modelled to be from P as well. In this case, the conditional null depends on random data, leading to severe difficulties in the interpretation of type one error. For this reason, in this paper we will only consider the unconditional error rate. However, setting $g = 0$ and plugging in a ready-made classification rule for Φ , regardless of the data it was learnt on, leads to a sort of conditional null hypothesis. In this case, the true error becomes a random variable of the learning data, and the latter *must not* be from the sample (z_1, \dots, z_n) . We will not go into details of conditional testing or of the case $g = 0$.

The form of H^0 suggests a t -test. However, the number of independent realizations of $\Phi(\mathbf{z}_1, \dots, \mathbf{z}_{g+1})$ is only $\lfloor n/(g+1) \rfloor$, since it is to be computed with respect to $P^{\otimes(g+1)}$. Therefore, a correct t -test would be severely underpowered, and cross-validation procedures are usually preferred.

2.4. Cross-validation. Let us now show how cross-validation procedures fit into the framework described above. In a cross-validation procedure, dependent realizations of $\Phi(\mathbf{z}_1, \dots, \mathbf{z}_{g+1})$ are considered. More precisely, for every ordered subset $T = (i_1, \dots, i_g; i_{g+1})$ of $\{1, \dots, n\}$,

$$(7) \quad \tilde{\Delta}(T) := \Phi(\mathbf{z}_{i_1}, \dots, \mathbf{z}_{i_g}; \mathbf{z}_{i_{g+1}})$$

is an estimator of Δ , where we visually separate learning and test sets again. We view $\tilde{\Delta}(T)$ as an estimator of the difference of error rates of classifiers learnt on samples of size g instead of n , in contrast to differing usage in the literature. Thus, $\tilde{\Delta}(T)$ is unbiased. Of course, the word ‘‘ordered subset’’ refers to the order $1 < \dots < g + 1$; it is not imposed that $i_1 < \dots < i_{g+1}$. Similarly, if $S = \{i_1, \dots, i_{g+1}\}$ is an *unordered* subset of size $g + 1$ of $\{1, \dots, n\}$, the value $\Phi_0(\mathbf{z}_{i_1}, \dots, \mathbf{z}_{i_{g+1}})$, using the symmetric Φ_0 instead of Φ , does not depend on the order of S . Therefore, we can unambiguously extend definition (7) to a function $\tilde{\Delta}$ also on the collection of unordered subsets S by setting $\tilde{\Delta}(S) := \Phi_0(S)$. This is an unbiased estimator of Δ . Also, let \mathcal{T} be a collection of ordered subsets T as above. Then, let

$$(8) \quad \tilde{\Delta}(\mathcal{T}) := \frac{1}{|\mathcal{T}|} \sum_{T \in \mathcal{T}} \tilde{\Delta}(T)$$

be the average of all values of $\tilde{\Delta}(T)$ over \mathcal{T} , and similarly $\tilde{\Delta}(\mathcal{S})$ for a collection \mathcal{S} of unordered subsets S as above the average of all values $\tilde{\Delta}(S)$ involving the symmetric function Φ_0 . Equation (8) may involve each learning set multiple times because each observation of a test sample can then contribute a summand to (8). In other contexts the mean error rate over the entire test sample is counted as only one occurrence of the learning set.

For any such collections \mathcal{T} or \mathcal{S} , the estimators $\tilde{\Delta}(T)$ and $\tilde{\Delta}(S)$ are unbiased for Δ . As soon as \mathcal{T} contains together with an ordered subset $T = (i_1, \dots, i_g; i_{g+1})$ all its cyclic permutations $(i_2, \dots, i_{g+1}; i_1)$, $(i_3, \dots, i_{g+1}, i_1; i_2)$ and so on, we have $\tilde{\Delta}(\mathcal{T}) = \tilde{\Delta}(\mathcal{S})$ where the collection \mathcal{S} is obtained from the collection \mathcal{T} by forgetting the order (and the multiple entries with the same order coming from the cyclic permutation).

Now, the ordinary K -fold cross-validation can be incorporated in this framework as follows. Suppose that K is such that $K(n - g) = n$, possibly after disregarding a few observations in order to assure divisibility of n by $n - g$. Therefore, $g \geq n/2$. The extreme cases are $g = n/2$ for $K = 2$ and $g = n - 1$ for $K = n$. Let \mathcal{T}_{CV} be a collection of ordered subsets of the form

$$(9) \quad T = (1, \dots, k(n - g), (k + 1)(n - g) + 1, \dots, n; t)$$

where $k = 0, \dots, K - 1$ enumerates the learning blocks, the notation is to be read in such a way that if $k = 0$ the first entry is $n - g + 1$ and if $k = K - 1$ the last one is n , and $t \in \{k(n - g) + 1, \dots, (k + 1)(n - g)\}$ enumerates all test observations distinct from the learning block. Thus, T consists of one or two learning strides whose indices are contiguous and whose sizes add up to g , together with a single test observation index distinct from any learning observation index. Then $\tilde{\Delta}(\mathcal{T}_{CV})$ recovers the ordinary cross-validation estimator of Δ . In practise, one may also compute $\tilde{\Delta}(\mathcal{T}_{CV})$ from a permutation of the data, but this does not influence the formal description because $P^{\otimes n}$ is permutation-invariant.

Definition 2.4. We will in general refer to estimators of the form $\tilde{\Delta}(\mathcal{T})$ or $\tilde{\Delta}(\mathcal{S})$ given by (8), as to *cross-validation-like procedures*, irrespectively of the structure of \mathcal{T} or \mathcal{S} .

It was shown in Bengio & Grandvalet (2003) that there is no unbiased estimator of the variance $\mathbb{V}_{P^{\otimes n}}(\tilde{\Delta}(\mathcal{T}_{CV}))$ for any cross-validation procedure \mathcal{T}_{CV} , i.e. any divisor K of n .

It seems plausible from this tedious description of cross-validation that such a particular design \mathcal{T}_{CV} consisting merely of sets of the special form (9) does not lead to a globally small variance of $\tilde{\Delta}(\mathcal{T}_{CV})$ among all possible designs \mathcal{T} with fixed learning set size g . This variance is minimal for the cross-validation-like procedure \mathcal{T}_{max} consisting of *all* size $(g + 1)$ -subsets. We will expose the cases where there is an unbiased variance estimator of it, in contrast to the cross-validation case. Let us call this \mathcal{T}_{max} the maximal design. Another immediate advantage of it over an incomplete one is the fact that the need for a balanced data set, i.e. algorithms whose class labels are equally frequent, and/or for balanced blocks falls away. The only case of g and K such that $\mathcal{T}_{CV} = \mathcal{T}_{max}$ is the leave-one-out case $g = n - 1$ respectively $K = n$.

Similarly, one can distinguish those cases of g respectively K such that the associated design \mathcal{T} contains along with an ordered subset T all its cyclic permutations. In such a case, $\tilde{\Delta}(\mathcal{T}) = \tilde{\Delta}(\mathcal{S})$ for the design \mathcal{S} corresponding to \mathcal{T} . Among the cross-validation procedures, only the leave-one-out case $g = n - 1$ respectively $K = n$ produces this situation. However, among the cross-validation-like procedures, this can happen for any g . For instance, it holds for the maximal design for all $0 \leq g \leq n - 1$. This is important to keep in mind for numerical implementation.

2.5. The full cross-validation-like estimator of Δ is a U -statistic. In this section, we show that the cross-validation-like procedure with maximal design, where all size- g -subsets of the sample are used for learning, is a U -statistic and therefore has least variance among all cross-validation-like procedures. It seems that this fact has not yet been described in the literature. Among the immediate consequences of interpreting this procedure as a U -statistic will be asymptotic normality, the first case of Theorem 4.1. The parameter of interest $\Theta = \Delta$ is a regular parameter because of (4); this equation also shows that its degree is at most $g + 1$.

Assumption 2.5. The degree of Δ is exactly $g + 1$.

This states that the true error rate cannot be computed from learning samples of smaller size than g for all (reasonable) distributions P . While it is not automatic, it seems to be violated only in irrelevant artificial counterexamples, such as for instance one of the form $\Phi(z^1, z^2, z^3) = \Phi'(z^2, z^3)$ where the learning step only makes use of a part of the learning set observations, and in similar cases. So, the assumption is natural.

Let \mathcal{S}_{max} and \mathcal{T}_{max} be the maximal designs of unordered and ordered subsets, respectively, as introduced above. The corresponding error rate estimator is then the U -statistic associated to the particular kernel $\Psi = \Phi$ and $\Psi_0 = \Phi_0$, respectively. We define

$$(10) \quad \hat{\Delta} := U(\Phi_0) = \Phi_0(\mathcal{S}_{max}) = \Phi(\mathcal{T}_{max})$$

as the associated U -statistic as in Definition 2.1, i.e. the one defined by the symmetric kernel Φ_0 . It follows immediately from Hoeffding (1948) that it has minimal variance among all unbiased estimators of Δ . In particular, it has strictly smaller variance than all cross-validation procedures for $2 \leq g \leq n - 2$, and is equal to the cross-validation estimator in the leave-one-out case $g = n - 1$. Lee (1990, Section 4.3, Theorems 1 and 4) describes quantitatively the variance decrease of $\hat{\Delta}$ with respect to $\tilde{\Delta}$. These theorems treat the case of a fixed \mathcal{S} and an \mathcal{T} consisting of random subsets, respectively. The statistic $\hat{\Delta}$ coincides with what is called complete cross-validation in Kohavi (1995), as well as with

complete repeated sub-sampling as considered in Boulesteix et al. (2008), or leave- p -out cross-validation in Shao (1993) and Arlot & Celisse (2010), where $p = n - g$.

In practise, the definition of $\widehat{\Delta}$ involves too many summands for computation, but can be easily approximated to arbitrary precision using an \mathcal{S} of random subsets, see Section 5.

3. A U -STATISTIC ESTIMATOR OF $\mathbb{V}(\widehat{\Delta})$

3.1. Variances are regular parameters. The theory of U -statistics comes to full power as soon as not only the original regular parameter Δ is estimated optimally by a U -statistic $\widehat{\Delta}$, but also the variance $\mathbb{V}(\widehat{\Delta})$ of this U -statistic itself is exhibited as another regular parameter, this time depending not only on Φ_0 but also on n . Therefore, we are in a position to estimate $\mathbb{V}(\widehat{\Delta})$ by a U -statistic as well.

In the following Proposition, we outline formally that variances and covariances are regular parameters in general, without determining optimally the degree. Thus, the full power of U -statistics can be used to estimate them. We then pin down the degree in Proposition 3.2.

Proposition 3.1. *Let $f(\mathbf{z}_1, \dots, \mathbf{z}_k)$ be a function of k realizations of independent identically distributed random variables $\mathbf{Z}_i \sim P$ with existing variance $\mathbb{V}_{P^{\otimes k}}(f) < \infty$. Then the variance $\mathbb{V}_{P^{\otimes k}}(f)$ is a regular parameter of degree at most $2k$. More generally, the covariance between two such functions f and g , as soon as it exists, is a regular parameter of degree at most $2k$.*

Proof. Both $\mathbb{V}(f)$ and $\text{cov}(f, g)$ are, by definition, polynomials of integrals with respect to P . In order to show that they are regular parameters, we have to rewrite each one as a single integral instead. This is accomplished by

$$(11) \quad \begin{aligned} \mathbb{V}_{P^{\otimes k}}(f) &= \mathbb{E}(f^2) - \mathbb{E}(f)^2 \\ &= \int \cdots \int \frac{1}{2} (f(\mathbf{z}_1, \dots, \mathbf{z}_k) - f(\mathbf{z}_{k+1}, \dots, \mathbf{z}_{2k}))^2 dP(\mathbf{z}_1) \cdots dP(\mathbf{z}_{2k}) \end{aligned}$$

and an almost analogous formula for the covariance $\text{cov}_{P^{\otimes k}}(f, g)$. \square

The integrand is not unique. It was chosen in such a way to resemble the symmetric kernel $(z_1 - z_2)^2/2$ of the variance of P itself, i.e. the case $r = 1, f(z_1) = z_1$. Furthermore, the degree of $\mathbb{V}(f)$, i.e. the minimal length of an integrand that accomplishes this, can be much smaller than $2k$ and depends on f . Also, the integrand of (11) is not symmetric in general and remains to be symmetrized.

Let us now investigate the case where f is a U -statistic associated to a symmetric kernel Φ_0 . Caution has to be taken because the regular parameter now depends on n , in sharp contrast to the U -statistic $\widehat{\Delta}$ itself. For the case $f = \widehat{\Delta}$, we have $k = n$, so our knowledge attained so far on the degree of the kernel of $\mathbb{V}(\widehat{\Delta})$ is that it is at most $2n$. However, it is possible to obtain better insight into the degree of the variance. It will turn out that the variance is a linear combination of regular parameters, each of whose degrees do *not* depend on n , only the coefficients of the linear combination depend on n . This is the content of the following proposition, which presents in short form results of Hoeffding (1948, Section 5) as well as immediate consequences.

In the following, we will consider a general underlying U -statistic U which estimates an unknown parameter Θ , and develop the theory of its variance as it is needed for its estimation. From Section 3.2 on, we will pay particular attention to the case where U is associated to the kernel Φ_0 defined by (5), thus $\Theta = \Delta$ and $U = \widehat{\Delta}$.

Proposition 3.2. *Let U be the U -statistic associated to a bounded symmetric kernel Φ_0 of degree m and to a total sample size n . Denote $\Theta = \mathbb{E}(\Phi_0)$. Then the variance of U is a regular parameter of degree at most $2m$. Furthermore, it splits as a sum*

$$(12) \quad \mathbb{V}(U) = \sum_{c=1}^m \alpha_c \kappa_c - (1 - \alpha_0) \Theta^2,$$

where α_c is the mass function at c of the hyper-geometric distribution $\mathcal{H}(n, m, m)$, and all κ_c are regular parameters satisfying

$$(13) \quad \kappa_c = \int \cdots \int \Phi_0(\mathbf{z}_1, \dots, \mathbf{z}_m) \Phi_0(\mathbf{z}_{m-c+1}, \dots, \mathbf{z}_{2m-c}) dP(\mathbf{z}_1) \cdots dP(\mathbf{z}_{2m-c}).$$

Thus, κ_c is a regular parameter of degree at most $2m - c$. Furthermore, since

$$(14) \quad \Theta^2 = \int \cdots \int \Phi_0(\mathbf{z}_1, \dots, \mathbf{z}_m) \Phi_0(\mathbf{z}_{m+1}, \dots, \mathbf{z}_{2m}) dP(\mathbf{z}_1) \cdots dP(\mathbf{z}_{2m}),$$

Θ^2 is a regular parameter of degree at most $2m$.

Proof. Direct computation shows that the right hand side of (13) coincides with what is called $\mathbb{E}(\Phi_c^2(X_1, \dots, X_c))$ in Hoeffding (1948, Section 5) for all $1 \leq c \leq m$. This step involves the symmetry of the kernel Φ_0 and careful renaming of the variables. Hoeffding already supposes a symmetric kernel which he calls Φ . The quantities ζ_c of Hoeffding (1948) – which are called σ_c in Lee (1990) – are thus related to our κ_c by means of the equation $\zeta_c = \kappa_c - \Theta^2$, as follows from Hoeffding (1948, formula 5.10). From $\mathbb{V}(U) = \sum_{c=1}^m \alpha_c \zeta_c$ (Hoeffding, 1948, 5.13) we thus deduce $\mathbb{V}(U) = \sum_{c=1}^m \alpha_c (\kappa_c - \Theta^2) = \sum_{c=1}^m \alpha_c \kappa_c - (1 - \alpha_0) \Theta^2$ because $\sum_{c=1}^m \alpha_c = (1 - \alpha_0)$.

The fact that linear combinations of regular parameters are regular parameters (Hoeffding, 1948, Page 295) completes the proof. \square

Proposition 3.2 achieves the desired simplification: The degree of $\mathbb{V}(U)$ for a U -statistic U of degree m is shown to be at most $2m$ instead of $2n$, and the dependence of $\mathbb{V}(U)$ on n is now expressed solely by means of the hyper-geometric mass function, whereas κ_c and $\mathbb{E}(U)^2$ do not depend on n .

Remark 3.3. Direct computation yields $\mathbb{E}(U^2) = \sum_{c=1}^m \alpha_c \kappa_c + \alpha_0 \Theta^2$, making use of the fact that the kernel is symmetrized. This together with the usual decomposition $\mathbb{V}(U) = \mathbb{E}(U^2) - \mathbb{E}(U)^2 = \mathbb{E}(U^2) - \Theta^2$ also proves (12) and shows that the degree of $\mathbb{E}(U^2)$ is at most $2m$. It is natural to assume that its degree is exactly $2m$, in analogy to assumptions 2.5 above and 3.4 below. In contrast, the advantage of decomposition (12) is that the first summand only involves the κ_c which all have smaller degree than $2m$, namely $2m - c$. Therefore, we prefer decomposition (12) over the usual decomposition and work with and estimate the quantities κ_c rather than Hoeffding's ζ_c which all have degree $2m$ (see also Remark 3.8 below).

3.2. Definition of the U -statistic for the variance. In order to estimate the variance of the U -statistic U by another U -statistic, we need the following.

Assumption 3.4. In the general situation of Proposition 3.2, the statistic U is non-degenerate. In the particular case $U = \widehat{\Delta}$ where $\Theta = \Delta$, this states that $\kappa_c \neq \Delta^2$ for $1 \leq c \leq g + 1$.

Furthermore, we assume in the situation of Proposition 3.2 that all upper bounds for the degrees thus obtained are optimal. In the particular case $U = \widehat{\Delta}$, this means that the degree of κ_c is $2m - c = 2g + 2 - c$ and that of $\Theta^2 = \Delta^2$ is $2m = 2g + 2$.

The non-degeneracy can be numerically checked for plausability, unlike Assumption 2.5 and the degree optimality which both state that the regular parameters cannot be written by a smaller number of integrals. There seems to be no reason why a kernel of the form (5) with a non-trivial classifier should not satisfy them. The first part of Assumption 3.4 is needed for the central limit theorem 4.1, the second one for Theorem 3.9.

Proposition 3.2 motivates the following definition.

Definition 3.5. In the general situation of Proposition 3.2 and under Assumption 3.4, the statistics $\widehat{\kappa}_c$ for $1 \leq c \leq m$ of degree $2m - c$ and the statistic $\widehat{\Theta}^2$ of degree $2m$ are defined as the U -statistics associated to the symmetrized versions of the kernels which are the integrands in (13) and (14), respectively.

Estimating Θ^2 by U^2 instead would be biased and thus would not fit in our framework.

Remark 3.6. It would not be suitable to simply estimate Θ^2 by zero in view of $H^0 : \Theta = 0$. The first reason is that failure to subtract $(1 - \alpha_0) \cdot \widehat{\Theta}^2$ from the variance estimator (15) below would overestimate the variance $\mathbb{V}(U)$ under H^1 , leading to a severe loss of power. The second is that it would conflict with Hoeffding's setup. In fact, under the null hypothesis $\Theta = 0$, the degree of Θ and that of Θ^2 would be trivially zero, if we were willing to restrict Hoeffding's class \mathcal{P} of distributions to only ones obeying the null; however, the least-variance optimality property of a U -statistic relies on \mathcal{P} encompassing *all* properly discontinuous distribution functions, not only null ones. Likewise, the degree has to be defined for a global class of null and alternative together. This is akin of a classical one-way analysis of variance statistic where estimating variance within and between groups separately greatly increases the power.

We can now define the variance estimator of a U -statistic as a U -statistic itself.

Definition 3.7. In the general situation and notation of Proposition 3.2 and under Assumption 3.4, we define an estimator, abbreviated \widehat{w} , for the variance of the U -statistic U as the U -statistic associated to the linear combination as in (12) of the kernels of κ_c and of Θ^2 given by (13) and (14).

After a short and straightforward computation, the definition can be re-stated alternatively in more explicit form: The single U -statistic \widehat{w} splits as a sum

$$(15) \quad \widehat{w} := \sum_{c=1}^m \alpha_c \widehat{\kappa}_c - (1 - \alpha_0) \widehat{\Theta}^2$$

of U -statistics of varying degrees. In the particular case $U = \widehat{\Delta}$ where $\Theta = \Delta$, this defines an estimator for $\mathbb{V}(\widehat{\Delta})$, which will be abbreviated by \widehat{v} .

The estimator \widehat{w} enjoys the unbiasedness and optimality properties analogous to $\widehat{\Delta}$. In particular, this applies to \widehat{v} .

Remark 3.8. In the latter case $U = \widehat{\Delta}$, we have $m = g + 1$ so the degree of \widehat{v} is $2g + 2$, and that of $\widehat{\kappa}_c$ was given in the degree optimality statement of Assumption 3.4. The reason for splitting \widehat{w} into several U -statistics of varying degree is numerical efficiency: Hoeffding (1948) suggests to estimate $\zeta_c = \kappa_c - \Theta^2$. However, all of these parameters have degree $2m$. Instead, it is of course advisable to estimate the κ_c which have smaller degree $2m - c$. Then, Θ^2 needs to be estimated only once. This remark is the empirical analogue to Remark 3.3.

3.3. Existence criterion and order of consistency. We are now in a position to investigate the estimator of $\mathbb{V}(\widehat{\Delta})$. In principle, this section applies to the general situation of Proposition 3.2, but in order to keep the presentation clear we will focus on the interesting case $\Theta = \Delta, U = \widehat{\Delta}$ for the rest of the paper. Therefore, we will write $\widehat{\Delta}^2$ for the statistic $\widehat{\Theta}^2$ whereas we will not introduce a special notation for the statistics κ_c for that case.

In the consistency statement of Theorem 3.9 below, the true parameter $\mathbb{V}(\widehat{\Delta})$ depends on n , unlike in a typical consistency statement. In principle, the sample size used for this estimation does not need to be the same n again, but can in fact be any number $n' \geq 2g + 2$. However, in practise the same sample is used to estimate Δ as well as $\mathbb{V}(\widehat{\Delta})$, so we restrict our attention to the diagonal case $n = n'$ for simplicity. This is analogous to the ordinary one-sample t -test statistic, where both the numerator, the sample mean, and its standard deviation, the denominator, are simultaneously estimated on the same sample, so with the same n . However, in our case, no factor $n^{-1/2}$ cancels out between both.

Theorem 3.9. *If $n \geq 2g + 2$, the estimator \hat{v} of $\mathbb{V}(\widehat{\Delta})$ has least variance among all unbiased estimators of $\mathbb{V}(\widehat{\Delta})$ over any family of distributions \mathcal{P} containing all purely discontinuous distribution functions. Furthermore, \hat{v} is strongly consistent in the sense that $n^{d/2}(\hat{v} - \mathbb{V}(\widehat{\Delta})) \rightarrow 0$ almost surely for any $0 \leq d \leq 2$.*

We do not claim to have exhibited the optimal order of consistency.

(of Theorem 3.9). The unbiasedness of \hat{v} as well as its least-variance optimality follow from the general properties of U -statistics. Only the consistency statement remains to be shown. For $0 \leq d < 1$, Hoeffding (1963, Equation 5.7) applied to the U -statistic \hat{v} whose kernel is bounded between 0 and 1, yields the quantitative version

$$(16) \quad P\left(\left|\hat{v} - \mathbb{V}(\widehat{\Delta})\right| \geq \varepsilon n^{-d/2}\right) \leq 2 \exp\left(-2\lfloor n/(2g+2) \rfloor \varepsilon^2 n^{-d}\right)$$

for any $\varepsilon > 0$ which has to be applied with care because the degree of the U -statistic varies with n . This only applies to $0 \leq d \leq 1$ and only shows weak consistency. In the following, we make use of the fact that U -statistics are strongly consistent meaning that they satisfy the strong law of large numbers if the kernel is absolutely integrable, for instance bounded. This was first shown in an unpublished paper of 1961 by Hoeffding, a complete proof is given in Lee (1990, Section 3.4.2, Theorem 3). For all cases $0 \leq d \leq 2$, let us first show that $n\hat{v}$ almost surely tends to $(g+1)^2(\kappa_1 - \Delta^2)$. For $c \geq 2$, the summands $n\alpha_c \widehat{\kappa}_c$ of $n\hat{v}$ almost surely tend to zero, because $n\alpha_c$ does, and $\widehat{\kappa}_c$ is strongly consistent, so the sequence $\widehat{\kappa}_c$ for $n \rightarrow \infty$ is almost surely bounded, for every c . For $c = 1$, the summand $n\alpha_1 \widehat{\kappa}_1$ almost surely tends to $(g+1)^2 \kappa_1$, because $n\alpha_1 \rightarrow (g+1)^2$, and $\widehat{\kappa}_1$ is strongly consistent. Similarly, the summand $n(1 - \alpha_0) \widehat{\Delta}^2$ almost surely tends to $(g+1)^2 \Delta^2$, because $n(1 - \alpha_0) \rightarrow (g+1)^2$, and $\widehat{\Delta}^2$ is strongly consistent.

The statement now follows from the fact that $\lim_{n \rightarrow \infty} n\mathbb{V}(\widehat{\Delta}) = (g+1)^2(\kappa_1 - \Delta^2)$ (Hoeffding, 1948, 5.23). \square

Under H^0 , there are unbiased estimators already for smaller m since then Δ^2 does not need to be estimated. However, as noted in Remark 3.6, the optimality property cannot be shown in this case.

4. TESTING

4.1. Central limit theorem. The convergence of $\widehat{\Delta}$ towards Δ as $n \rightarrow \infty$ is described by the Strong Law of Large Number, the Law of the Iterated Logarithm and the Berry-Esseen

theorem (Lee, 1990, Section 3.4.2, Theorem 3, Section 3.5, Theorem 1, Section 3.3.2, Theorem 1, respectively). In order to show the existence of an asymptotically exact test, we need the following theorem as it subsumes the unstudentized and the studentized case. It is reminiscent of and contains as special case the statement that the t -distributions tend to $\mathcal{N}(0, 1)$ as the degrees of freedom tend to infinity.

Theorem 4.1. *Let $u(n)$ be one of the following expressions: the asymptotic variance $u(n) := (g + 1)^2(\kappa_1 - \Delta^2)/n$, the expression $u(n) := (g + 1)^2(\widehat{\kappa}_1 - \widehat{\Delta}^2)/n$, where $\widehat{\kappa}_1$ and $\widehat{\Delta}^2$ are defined by the case $\Theta = \Delta, U = \widehat{\Delta}$ in Definition 3.5, or $u(n) := \widehat{v}$ as of Definition 3.7. Then $(\widehat{\Delta} - \Delta)u(n)^{-1/2}$ converges in distribution to $\mathcal{N}(0, 1)$ as g remains fixed, $n \rightarrow \infty$.*

The occurrence of the factor $(g + 1)^2$ is explained by the fact that this is the decay rate of the coefficient α_1 in the sense that $\lim_{n \rightarrow \infty} n\alpha_1 = (g + 1)^2$. This also explains the asymptotic behaviour of the variance.

The first case of Theorem 4.1 shows approximate normality of the unstudentized statistic $\widehat{\Delta}$ itself. It seems that there exists no statement in the literature giving the precise reason why a cross-validation type estimator is asymptotically normally distributed. This case appears in the asymptotic variance statement Hoeffding (1948, 5.23). The second case is included for systematic reasons; this expression is the empirical analogue of the first case, but is a biased variance estimator. Finally, the third case includes the unbiased variance estimator elaborated in the present manuscript. The fact that $\mathbb{V}(\widehat{\Delta})/u(n)$ tends to one, shown in the following proof, is not immediate, due to the diagonality property $n = n'$ mentioned above. Likewise, it is not obvious whether this ratio almost surely tends to one.

(of Theorem 4.1). In the first case, this is Hoeffding (1948, Theorem 7.1) and rests on the validity of the first part of Assumption 3.4. In the other cases, the proof proceeds simultaneously. First, the proof of Theorem 3.9 shows that convergence of $n\mathbb{V}(\widehat{\Delta})$ implies the almost sure convergence not only of \widehat{v} but in fact of $nu(n)$ for any of the choices of $u(n)$. Thus, $(nu(n))^{-1}$ is almost surely bounded. This statement is licit because the first part of Assumption 3.4 implies $\kappa_1 \neq \Delta^2$, so $nu(n)$ converges to a non-zero value and, therefore, there are at most only finitely many n such that $u(n) = 0$ has positive probability. Consequently, we may multiply $n(u(n) - \mathbb{V}(\widehat{\Delta}))$ which converges almost surely to zero, hence also in probability, with $(nu(n))^{-1}$ to show that in each case, the ratio $\mathbb{V}(\widehat{\Delta})/u(n)$ tends to one in probability by Slutsky's theorem. By the continuous mapping theorem, $(\mathbb{V}(\widehat{\Delta})/u(n))^{1/2}$ tends to one in probability as well. Therefore, $(\widehat{\Delta} - \Delta)u(n)^{-1/2} = (\widehat{\Delta} - \Delta)(\mathbb{V}(\widehat{\Delta}))^{-1/2}(\mathbb{V}(\widehat{\Delta})/u(n))^{1/2}$ tends to $\mathcal{N}(0, 1)$ in distribution by the first case and another application of Slutsky's theorem. \square

4.2. Asymptotic rejection regions and confidence intervals. So, the two-sided test of H^0 with the rejection region

$$(17) \quad \left\{ \left| \widehat{\Delta} \right| \geq u(n)^{1/2} \phi^{-1}(1 - \alpha/2) \right\}$$

has asymptotic level α , where ϕ is the standard normal cumulative distribution function. While the second case of Theorem 4.1 uses a positively biased variance estimator and hence provides a conservative test which, however, is asymptotically exact, the third case provides the best approximation to exactness already in the finite case. Likewise, an asymptotically exact confidence interval for Δ at level $1 - \alpha$ is

$$(18) \quad \left[\widehat{\Delta} - u(n)^{1/2} \phi^{-1}(1 - \alpha/2), \widehat{\Delta} + u(n)^{1/2} \phi^{-1}(1 - \alpha/2) \right].$$

A related, but different approach to a similar testing problem is provided by the so-called empirical Bernstein inequalities in Peel et al. (2010, Equations 12, 13). These are sharp empirical inequalities for general U -statistics associated to bounded kernels. However, n has to be an integer multiple of the degree, and the authors do not consider cross-validation, but only partitions of the test set.

5. THE CONVERGENCE OF INCOMPLETE TO COMPLETE U -STATISTICS IN PRACTISE

In practical applications, the number of summands of (1) is too large for computation. In the particular case where one of the learning methods M is a k -nearest-neighbour algorithm, it is possible to compute the corresponding summand of the complete U -statistic, the leave- p -out cross-validation estimator of the error rate, by an efficient closed-form expression (Celisse & Mary-Huard, 2012). In general, however, one can only consider a design \mathcal{T} smaller than the full \mathcal{T}_{max} , leading to incomplete U -statistics as treated in Lee (1990), for instance. We now show that the incomplete U -statistic with random design approximates the complete one satisfactorily after a feasible number of iterations.

Let Φ be a not necessarily symmetric kernel with $-1 \leq \Phi \leq 1$, let \mathcal{T}^* be a collection of N randomly drawn ordered size m -subsets of $\{1, \dots, n\}$ from the equidistribution Q on the collection of such subsets, and let $\Phi(\mathcal{T}^*)$ be the associated incomplete U -statistic. Then the probability of approximation error at least $\delta > 0$ is bounded by

$$(19) \quad pr_Q(|\Phi(\mathcal{T}^*) - \Phi(\mathcal{T})| \geq \delta) \leq 2 \exp(-\delta^2 N/2).$$

This follows from Hoeffding (1963, Theorem 2) because the entries of \mathcal{T}^* were drawn independently from each other. One should be aware that here we do *not* refer to the part of Hoeffding (1963) concerned with U -statistics, in contrast to the situation of the similar inequality (16), where we did so. Here, we formulated the version for ordered subsets because this is of immediate interest for computation.

The fast exponential decay of (19) implies that sufficiency of the approximation is assured as soon as N is a small multiple of δ^{-2} , where δ is the pre-specified tolerance. Precisely, the following corollary of Hoeffding's theorem can be used in practise.

Corollary 5.1. *After at most $N = 2d + 1$ iterations, d digits after the comma are fixed with a probability of at least $1 - 2 \exp(-5) \approx 0.99$.*

Such a number of repetitions is, in general, hard but feasible because this N is the mere number of times a model has to be fitted. For instance, in the illustration in Section 6, no tuning of the hyper-parameter λ is part of each iteration. Remarkably, this bound on N holds true irrespectively of the sample size n or of any properties of the particular U -statistic under consideration, apart from $-1 \leq \Phi \leq 1$. In practise, however, one proceeds again slightly differently. For the case of approximation of the U -statistic $\hat{\Delta}$ for instance, one applies the following procedure which yields even faster convergence against the true $\hat{\Delta}$. In the formal setting required for inequality (19), one would use only one test observation for each learning iteration, which would lead to unnecessarily high computational cost. Instead, one simply uses all remaining $n - g$ observations for testing. This speeds up convergence even further. Corollary 5.1 also applies to the computation of \hat{v} by the linear combination of U -statistics explained above because the kernels appearing in (13) and (14) are bounded between -1 and 1 as well.

6. THE CALCULATIONS IN A REAL DATA EXAMPLE

The estimation procedure elaborated in the preceding sections was applied to the well-investigated colon cancer data set by Alon et al. (1999), where the binary response $y \in \mathcal{Y}$ stands for the type of tissue (either normal tissue or tumor tissue) and the 2000 continuous predictors are gene expressions. We used lasso-penalized logistic regression with the coordinate descent method for classification (Friedman et al., 2010) and the penalization parameters $\lambda = 0.08, 0.5$. Pre-selecting these values led to the software-internal estimator for the difference of error rates to be greater than 0.1. This involved the whole data set, however, this is no problem here. Sample size was $n = 62$. Therefore, the condition $n \geq 2g + 2$ constrained $g \leq 30$. Since the variance of the U -statistic \hat{v} decreases to the extent to which the sample size exceeds the degree $2g + 2$, the learning set size g was arbitrarily chosen to be only 26 to compromise with the effort to avoid a too small learning set size.

There were numerical evidence for the validity of the non-degeneracy statement of Assumption 3.4. The resulting point estimate of Δ was -0.14 , with 95%-confidence interval $[-0.35, 0.07]$ and estimated variance $\hat{v} = 0.01$. The number of iterations was $N = 10^5$ for each of the U -statistics $\hat{\Delta}, \hat{\kappa}_c$ for $1 \leq c \leq g + 1$ and $\hat{\Delta}^2$. By Corollary 5.1, two digits of each of these were therefore assured. The two-sided p -value was $p = 0.19$, given by the corresponding upper and lower normal tail probabilities. An R-script that allows to reproduce these results is available on the first author's institution web page.

ACKNOWLEDGEMENT

MF was supported by the German Science Foundation (DFG-Einzelförderung BO3139/2-2 to ALB). RH was supported by the German Science Foundation (DFG-Einzelförderung BO3139/3-1 to ALB). RDB was supported by the German Science Foundation (DFG-Einzelförderung BO3139/4-1 to ALB).

REFERENCES

- ALON, U., BARKAI, N., NOTTERMAN, D. A., GISH, K., YBARRA, S., MACK D. & LEVINE, A. J. (1999). Broad patterns of gene expression revealed by clustering analysis of tumor and normal colon tissue probed by oligonucleotide arrays. *Proc. Natl. Acad. Sci. USA* **96**, 6745–6750.
- ARLOT, S. & CELISSE, A. (2010). A survey of cross-validation procedures for model selection. *Stat. Surveys* **4**, 40–79.
- BENGIO, Y. & GRANDVALET, Y. (2003). No unbiased estimator of the variance of K-fold cross-validation. *J. Mach. Learn. Res.* **5**, 1089–1105.
- BOULESTEIX, A.-L., PORZELIUS, C. & DAUMER, M. (2008). Microarray-based classification and clinical predictors: on combined classifiers and additional predictive value. *Bioinformatics* **24**, 1698–1706.
- BRAGA-NETO, U. M. & DOUGHERTY, E. R. (2004). Is cross-validation valid for small-sample microarray classification? *Bioinformatics* **20**, 374–380.
- BÜHLMANN, P. & VAN DE GEER, S. (2011). *Statistics for High-Dimensional Data*. Springer Series in Statistics.
- CELISSE, A. & MARY-HUARD, T. (2012). Exact Cross-Validation for kNN: application to passive and active learning in classification. *J. Soc. Fr. Stat.* **152**, 83–97.
- DIETTERICH, T. G. (1998). Approximate statistical tests for comparing supervised classification learning algorithms. *Neural comput.* **10**, 1895–1923.

- DOUGHERTY, E. R., ZOLLANVARI, A. & BRAGA-NETO, U. M. (2011). The illusion of distribution-free small-sample classification in genomics. *Curr. genomics* **12**, 333.
- FRIEDMAN, J., HASTIE, T. & TIBSHIRANI, R. (2010). Regularization paths for generalized linear models via coordinate descent. *J. Stat. Softw.* **33**, 1–22.
- HOEFFDING, W. (1948). A class of statistics with asymptotically normal distribution. *Ann. Math. Stat.* **19**, 293–325.
- HOEFFDING, W. (1963). Probability inequalities for sums of bounded random variables. *J. Am. Statist. Assoc.* **58**, 13–30.
- JIANG, B., ZHANG, X. & CAI, T. (2008). Estimating the confidence interval for prediction errors of support vector machine classifiers. *J. Mach. Learn. Res.* **9**, 521–540.
- JIANG, W., VARMA, S. & SIMON, R. (2008). Calculating confidence intervals for prediction error in microarray classification using resampling. *Stat. Appl. Genet. Molec. Biol.*, **7**.
- KOHAVI, R. (1995). A study of cross-validation and bootstrap for accuracy estimation and model selection. *International Joint Conferences on Artificial Intelligence* **14**, 1137–1145.
- LEE, J. (1990). *U-statistics: Theory and Practice*. CRC Press.
- NADEAU, C. & BENGIO, Y. (2003). Inference for the generalization error. *Machine Learning* **52**, 239–281.
- PEEL, T., ANTHOINE, S. & RALAIVOLA, L. (2010). Empirical Bernstein inequalities for u-statistics. *Adv. Neural Inf. Process. Syst.* **23**, 1903–1911.
- SHAO, J. (1993). Linear model selection by cross-validation. *J. Am. Statist. Assoc.* **88**, 486–494.
- VAN DE WIEL, M., BERKHOF, J. & VAN WIERINGEN, W. (2009). Testing the prediction error difference between two predictors. *Biostatistics* **10**, 550–560.

INSTITUT FR MEDIZINISCHE INFORMATIONSVERARBEITUNG BIOMETRIE UND EPIDEMIOLOGIE, LUDWIG-MAXIMILIANS-UNIVERSITÄT MÜNCHEN, MARCHIONINISTR. 15, 81377 MÜNCHEN, GERMANY
E-mail address: {fuchs,hornung,debin,boulestex}@ibe.med.uni-muenchen.de