arXiv:2306.06140v1 [stat.ME] 9 Jun 2023

Null/No Information Rate (NIR): a statistical test to assess if a classification accuracy is *significant* for a given problem

M. Bicego, A. Mensi
University of Verona (Italy)

June 13, 2023

Abstract

In many research contexts, especially in the biomedical field, after studying and developing a classification system a natural question arises: "Is this accuracy *enough* high?", or better, "Can we say, with a statistically significant confidence, that our classification system is able to *solve* the problem"? To answer to this question we can use the statistical test described in this paper, which is referred in some cases as *NIR* (No Information Rate or Null Information Rate).

1 Motivation

In many research contexts, especially in the biomedical field, we have a classification problem for which we develop a classification system. Then we evaluate the performances of such system by measuring its *classification accuracy (or error)*, typically estimated with a Cross Validation protocol. In particular we have a dataset, which contains objects for which we know the true category, and we split the dataset in two separated sets: one, called *training set*, is used to build the classifier and the other, called *testing set*, is used to test it: we classify the objects in the testing set with the trained classifier and we count the number of times our classifier provides a correct answer, i.e. the answer of the classifier on a given object is identical to its true label.

At this point a simple question may arise: "Is this accuracy *enough* high?", or better, "Can we say, with a statistically significant confidence, that our classification system is able to *solve* the problem"? To answer to this question we can use the statistical test described here. In particular we aim at assessing, with a statistical test, if the computed classification accuracy is better than those which would have been obtained when:

- we assign every object of the testing set to a class which is *randomly chosen*

- we assign every object of the testing set to the class which is *more common* in the problem

The first represents just the accuracy of a *random* classifier (**Random**), whereas the second is typically referred to as **NIR** – No Information Rate or Null Information Rate. This second can also be considered as the result of the *a priori* classifier, a Bayesian classifier which assigns every object to the class ω_i which a priori probability $P(\omega_i)$ is maximum. $P(\omega_i)$ is estimated on the training set as:

$$P(\omega_i) = \frac{N_i}{N} \tag{1}$$

where N_i is the number of objects belonging to the i-th class, over the whole set of N objects in the training set [1]. While the intuition behind the first is trivial (just a random answer), for the second we exploit the a priori knowledge on the problem: if we know that one class is more frequent than another (for example the "control" class with respect to the "disease" class), then we can assign every testing object to such class. This represents another baseline, obtained with a very naive strategy which, at least, reflects the nature of the problem.

2 Notation

Before describing the test, let us introduce some notation. Given a dataset for a problem with C classes, let us call X^T and X^E the training and the testing sets, respectively, composed by the objects (x_i) and the corresponding true labels (y_i):

$$X^T = \{(x_1^T, y_1^T) \cdots (x_n^T, y_n^T)\}, \quad X^E = \{(x_1^E, y_1^E) \cdots (x_m^E, y_m^E)\}, \quad X^T \cap X^E = \emptyset$$

Within the Cross Validation protocol, we use the training set X^T to build the classifier $\mathcal{C}$; then we use the trained classifier to classify every object in the testing set X^E. The accuracy of the classifier $\mathcal{C}$ on X^E can be estimated by counting the number of correctly classified testing objects:

$$t(\mathcal{C}) = \sum_{i=1}^{m} \mathbf{I}(y_i^E = \mathcal{C}(x_i^E)) \tag{2}$$

where $\mathcal{C}(x_i^E)$ is the classification provided by the classifier $\mathcal{C}$ on the testing object x_i^E, and $\mathbf{I}(cond)$ is 1 if the condition "*cond*" is true, 0 otherwise. From t we can compute the accuracy as:

$$acc(\mathcal{C}) = \frac{t(\mathcal{C})}{m} \tag{3}$$

The baselines against which we want to compare are represented by the accuracy of the **random classifier** and the **NIR value**. The former is easily defined as:

$$accRand = \frac{1}{C} \tag{4}$$

while for the latter we have:

$$NIR = \frac{|X_a^E|}{m} \tag{5}$$

where X_a^E is the subset of X^E which contains only objects of class a, a represents the most frequent class inside the training set X^T and $|\ |$ represents the cardinality of a set.

3 The test

The test we describe here is also used in the R Caret Package [2], in particular to compare the observed accuracy to NIR. This package, and the test contained therein, has been widely applied in other works, such as [3]. The test is based on a one-sided binomial test: even though there is a lack of studies specifically related to this aspect, the use of a binomial test can be easily supported by considering a correctly classified sample as a success. Indeed, it is not uncommon to model (or think of) a classification task as a binomial experiment [4, 5] (paragraph 4.6.1), in which we can have a *success* or a *failure* depending on a correct or uncorrect classification. In detail, a Binomial test [6] evaluates whether the observed proportion of successes is significantly different from the expected proportion, which in our case corresponds to the successes of the Random classifier accuracy or the NIR. Please note that this is not an inferential test and it directly returns a p-value.

Given the classifier $\mathcal{C}$ and its corresponding accuracy on the testing set $acc(\mathcal{C})$, let us assume that $acc(\mathcal{C}) \geq NIR$ (or $acc(\mathcal{C}) \geq accRand$ if our baseline is the random accuracy) and that we want to check if the improvement is due to chance or not. To this end we perform the **binomial one-tailed test** as follows:

- We set the baseline against which we want to compare: in particular we define the expected probability of success p and the expected probability of failure q as :

  $$p = NIR, \qquad q = 1 - p = 1 - NIR$$

  (or $p = accRand, q = 1 - accRand$ for the test with respect to the random classifier).

- We set the number of observed successes as $t(\mathcal{C})$, i.e. the number of objects in the testing set X^E which have been correctly classified by $\mathcal{C}$.

- We compute the p-value $pval$, which represents the probability that the increase in the accuracy of $\mathcal{C}$ with respect to the NIR (or the Random accuracy) is due to chance, as:

  $$pval = \sum_{k=t(\mathcal{C})}^{m} \binom{m}{k} p^k q^{m-k} \tag{6}$$

This corresponds to a one-sided binomial test with parameters m, p and q

- If $pval \leq \alpha$ – where usually $\alpha \leq 0.05$ – then the classifier $\mathcal{C}$ is significantly better than the NIR classifier in terms of accuracy.

4 Final Remarks

- We can also compute the two-tailed p-value $pval2$, which is the probability that the difference between $acc(\mathcal{C})$ and the NIR (or the Random accuracy) is due to chance (without assuming that $acc(\mathcal{C}) \geq NIR$ or $acc(\mathcal{C}) \geq accRand$). In detail:

$$pval2 = 2 \sum_{k=t(\mathcal{C})}^{m} \binom{m}{k} p^k q^{m-k}. \tag{7}$$

- Please note that in [6] they state that whenever $mpq \geq 5$ the binomial distribution can be approximated to a normal one and a *z-statistic can be computed*. (No mention of this has been made in [2]).

- Matlab code is available from `https://profs.sci.univr.it/~bicego/code.html`.

References

[1] Richard O Duda, Peter E Hart, et al. *Pattern classification*. John Wiley & Sons, 2006.

[2] Max Kuhn. Building predictive models in r using the caret package. *Journal of Statistical Software*, 28(5):1–26, 2008.

[3] Adam Mourad Chekroud, Ryan Joseph Zotti, Zarrar Shehzad, Ralitza Gueorguieva, Marcia K Johnson, Madhukar H Trivedi, Tyrone D Cannon, John Harrison Krystal, and Philip Robert Corlett. Cross-trial prediction of treatment outcome in depression: a machine learning approach. *The Lancet Psychiatry*, 3(3):243–250, 2016.

[4] Pang-Ning Tan, Michael Steinbach, and Vipin Kumar. *Introduction to data mining*. Pearson Education India, 2016.

[5] Stan Aronoff et al. Classification accuracy: a user approach. *Photogrammetric Engineering and Remote Sensing*, 48(8):1299–1307, 1982.

[6] Bernard Rosner. *Fundamentals of biostatistics*. Cengage learning, 2015.