



Voting with random classifiers (VORACE): theoretical and experimental analysis

Cristina Cornelio¹ · Michele Donini² · Andrea Loreggia³ · Maria Silvia Pini⁴ · Francesca Rossi⁵

Accepted: 28 April 2021
© The Author(s) 2021

Abstract

In many machine learning scenarios, looking for the best classifier that fits a particular dataset can be very costly in terms of time and resources. Moreover, it can require deep knowledge of the specific domain. We propose a new technique which does not require profound expertise in the domain and avoids the commonly used strategy of hyper-parameter tuning and model selection. Our method is an innovative ensemble technique that uses voting rules over a set of randomly-generated classifiers. Given a new input sample, we interpret the output of each classifier as a ranking over the set of possible classes. We then aggregate these output rankings using a voting rule, which treats them as preferences over the classes. We show that our approach obtains good results compared to the state-of-the-art, both providing a theoretical analysis and an empirical evaluation of the approach on several datasets.

Keywords Multi-agent learning · Machine learning · Social choice theory

This work was mainly conducted prior joining Amazon.

✉ Andrea Loreggia
andrea.loreggia@gmail.com

Cristina Cornelio
cor@zurich.ibm.com

Michele Donini
donini@amazon.com

Maria Silvia Pini
pini@dei.unipd.it

Francesca Rossi
francesca.rossi2@ibm.com

¹ IBM Research, Rüschlikon, Zürich, Switzerland

² Amazon, Berlin, Germany

³ European University Institute, Firenze, Italy

⁴ Department of Information Engineering, University of Padova, Padova, Italy

⁵ IBM Research, Yorktown Heights, New York, USA

1 Introduction

It is not easy to identify the best classifier for a certain complex task [4, 25, 45]. Different classifiers may be able to exploit better the features of different regions of the domain at hand, and consequently their accuracy might be better only in that region [5, 29, 40]. Moreover, fine-tuning the classifier's hyper-parameters is a time-consuming task, which also requires a deep knowledge of the domain and a good expertise in tuning various kinds of classifiers. Indeed, the main approaches to identify the hyper-parameters' best values are either manual or based on grid search, although there are some approaches based on random search [6]. However, it has been shown that in many scenarios there is no single learning algorithm that can uniformly outperform the others over all data sets [22, 32, 46]. This observation led to an alternative approach to improve the performance of a classifier, which consists of combining several different classifiers (that is, an *ensemble* of them) and taking the class proposed by their combination. Over the years, many researchers have studied methods for constructing good ensembles of classifiers [16, 22, 30, 32, 42, 46], showing that indeed ensemble classifiers are often much more accurate than the individual classifiers within the ensemble [30]. Classifiers combination is widely applied to many different fields, such as urban environment classification [3, 53] and medical decision support [2, 49]. In many cases, the performance of an ensemble method cannot be easily formalized theoretically, but it can be easily evaluated on an experimental basis in specific working conditions (that is, a specific set of classifiers, training data, etc.).

In this paper we propose a new ensemble classifier method, called VORACE, which aggregates randomly generated classifiers using voting rules in order to provide an accurate prediction for a supervised classification task. Besides the good accuracy of the overall classifier, one of the main advantages of using VORACE is that it does not require specific knowledge of the domain or good expertise in fine-tuning the classifiers' parameters.

We interpret each classifier as a voter, whose vote is its prediction over the classes, and a voting rule aggregates such votes to identify the "winning" class, that is, the overall prediction of the ensemble classifier. This use of voting rules is within the framework of maximum likelihood estimators, where each vote (that is, a classifier's rank of all classes) is interpreted as a noisy perturbation of the correct ranking (that is not available), so a voting rule is a way to estimate this correct ranking [11, 13, 50].

To the best of our knowledge, this is the first attempt to combine randomly generated classifiers, to be aggregated in an ensemble method, using voting theory to solve a supervised learning task without exploiting any knowledge of the domain. We theoretically and experimentally show that the usage of generic classifiers in an ensemble environment can give results that are comparable with other state-of-the-art ensemble methods. Moreover, we provide a closed formula to compute the performance of our ensemble method in the case of Plurality, this corresponds to the probability of choosing the correct class, assuming that all the classifiers are independent and have the same accuracy. We then relax these assumptions by defining the probability of choosing the right class when the classifiers have different accuracies and they are not independent.

Properties of many voting rules have been studied extensively in the literature [24, 50]. So another advantage of using voting rules is that we can exploit that literature to make sure certain desirable properties of the resulting ensemble classifier hold. Besides the classical properties that the voting theory community has considered (like anonymity, monotonicity, IIA, etc.), there may be also other properties not yet considered, such as various forms of fairness [39, 47], whose study is facilitated by the use of voting rules.

The paper is organized as follows. In Sect. 2 we briefly describe some prerequisites (a brief introduction to ensemble methods and voting rules) necessary for what follows and an overview of previous works in this research area. In Sect. 3 we present our approach that exploits voting theory in the ensemble classifier domain using neural networks, decision trees, and support vector machines. In Sect. 4 we show our experimental results, while in Sects. 5, 6 and 7 we discuss our theoretical analysis: in Sect. 5 we present the case in which all the classifiers are independent and with the same accuracy, in Sect. 6 we relate our results with the Condorcet Jury Theorem also showing some interesting properties of our formulation (e.g. monotonicity and behaviour with infinite voters/classifiers); and in Sect. 7 we extend the results provided in Sect. 5 relaxing the assumptions of having all the classifiers with the same-accuracy and independent between each other. Finally, in Sect. 8 we summarize the results of the paper and we give some hints for future work.

A preliminary version of this work has been published as an extended abstract at the International Conference On Autonomous Agents and Multi-Agent Systems (AAMAS-20) [14]. The code is available open source at <https://github.com/aloreggia/vorace/>.

2 Background and related work

2.1 Ensemble methods

Ensemble methods combine multiple classifiers in order to give a substantial improvement in the prediction performance of learning algorithms, especially for datasets which present non-informative features [26]. Simple combinations have been studied from a theoretical point of view, and many different ensemble methods have been proposed [30]. Besides simple standard ensemble methods (such as averaging, blending, staking, etc.), *Bagging* and *Boosting* can be considered two of the main state-of-the-art ensemble techniques in the literature [46]. In particular, **Bagging** [7] trains the same learning algorithm on different subsets of the original training set. These different training subsets are generated by randomly drawing, with replacement, N instances, where N is the original size of training set. Original instances may be repeated or left out. This allows for the construction of several different classifiers where each classifier can have specific knowledge of part of the training set. Aggregating the predictions of the individual classifiers leads to the final overall prediction. Instead, **Boosting** [21] keeps track of the learning algorithm performance in order to focus the training attention on instances that have not been correctly learned yet. Instead of choosing training instances at random from a uniform distribution, it chooses them in a manner as to favor the instances for which the classifiers are predicting a wrong class. The final overall prediction is a weighted vote (proportional to the classifiers' training accuracy) of the predictions of the individual classifiers.

While the above are the two main approaches, other variants have been proposed, such as Wagging [54], MultiBoosting [54], and Output Coding [17]. We compare our work with the state-of-the-art in ensemble classifiers, in particular XGBoost [9], which is based on boosting, and Random Forest (RF) [27], which is based on bagging.

2.2 Voting rules

For the purpose of this paper, a *voting rule* is a procedure that allows a set of voters to collectively choose one among a set of candidates. Voters submit their vote, that is, their

preference ordering over the set of candidates, and the voting rule aggregates such votes to yield a final result (the winner). In our ensemble classification scenario, the voters are the individual classifiers and the candidates are the classes. A vote is a ranking of all the classes, provided by an individual classifier. In the classical voting setting, given a set of n voters (or agents) $A = \{a_1, \dots, a_n\}$ and m candidates $C = \{c_1, \dots, c_m\}$, a *profile* is a collection of n total orders over the set of candidates, one for each voter. So, formally, a voting rule is a map from a profile to a winning candidate.¹ The voting theory literature includes many voting rules, with different properties. In this paper, we focus on four of them, but the approach is applicable also to any other voting rules:

1. *Plurality* Each voter states who the preferred candidate is, without providing information about the other less preferred candidates. The winner is the candidate who is preferred by the largest number of voters.
2. *Borda* Given m candidates, each voter gives a ranking of all candidates. Each candidate receives a score for each voter, based on its position in the ranking: the i -th ranked candidate gets the score $m - i$. The candidate with the largest sum of all scores wins.
3. *Copeland Pairs* of candidates are compared in terms of how many voters prefer one or the other one, and the winner of such a pairwise comparison is the one with the largest number of preferences over the other one. The overall winner is the candidate who wins the most pairwise competitions against all the other candidates.
4. *Kemeny* [28] We borrow a formal definition of the rule from [12]. For any two candidates a and b , given a ranking r and a vote v , let $\delta_{a,b}(r, v) = 1$ if r and v agree on the relative ranking of a and b (e.g., they either both rank a higher, or both rank b higher), and 0 if they disagree. Let the agreement of a ranking r with a vote v be given by $\sum_{a,b} \delta_{a,b}(r, v)$, the total number of pairwise agreements. A Kemeny ranking r maximizes the sum of the agreements with the votes $\sum_v \sum_{a,b} \delta_{a,b}(r, v)$. This is called a Kemeny consensus. A candidate is a winner of a Kemeny election if it is the top candidate in the Kemeny consensus for that election.

It is easy to see that all the above voting rules associate a score to each candidate (although different voting rules associate different scores), and the candidate with the highest score is declared the winner. Ties can happen when more than one candidate results with the highest score, we arbitrarily break the tie lexicographically in the experiments. We plan to test the model on different and more fair tie-breaking rules. It is important to notice that when the number of candidates is $m = 2$ (that is, we have a binary classification task) all the voting rules have the same outcome since they all collapse to the **Majority** rule, which elects the candidate which has a majority, that is, more than half the votes.

Each of these rules has its advantages and drawbacks. Voting theory provides an axiomatic characterization of voting rules in terms of desirable properties such as anonymity, neutrality, etc. – for more details on voting rules see [1, 48, 50]. In this paper, we do not exploit these properties to choose the “best” voting rule, but rather we rely on what the experimental evaluation tells us about the accuracy of the ensemble classifier.

¹ We assume that there is always a unique winning candidate. In case of ties between candidates, we will use a predefined tie-breaking rule to choose one of them to be the winner.

2.3 Voting for ensemble methods

Preliminary techniques from voting theory have already been used to combine individual classifiers in order to improve the performance of some ensemble classifier methods [5, 18, 22, 31]. Our approach differs from these methods in the way classifiers are generated and how the outputs of the individual classifiers are aggregated. Although in this paper we report results only against recent bagging and boosting techniques of ensemble classifiers, we compared our approach with the other existing approaches as well. More advanced work has been done to study the use of a specific voting rule: the use of *majority* to ensemble a profile of classifiers has been investigated in the work of [34], where they theoretically analyzed the performance of majority voting (with *rejection* if the 50% of consensus is not reached) when the classifiers are assumed independent. In the work of [33], they provide upper and lower limits on the majority vote accuracy focusing on dependent classifiers. We perform a similar analysis of the dependence between classifier but in the more complex case of plurality, with also an overview of the general case. Although majority seems to be easier to evaluate compared to plurality, there have been some attempts to study plurality as well: [37] demonstrated some interesting theoretical results for independent classifiers, and [41] extended their work providing a theoretical analysis of the probability of electing the correct class by an ensemble using plurality, or plurality with rejection, as well as a stochastic analysis of the formula, and evaluating it on a dataset for human recognition. However, we have noted an issue with their proof: the authors assume independence between the random variable expressing the total number of votes received by the correct class and the one defining the maximum number of votes among all the wrong classes. This false assumption leads to a wrong final formula (the proof can be found in “Appendix” A). In our work, we provide a formula that exploits generating functions and that fixes the problem of [41], based on a different approach. Moreover, we provide proof for the two general cases in which the accuracy of the individual classifiers is not homogeneous, and where classifiers are not independent. Furthermore, our experimental analysis is more comprehensive: not limiting to plurality and considering many datasets of different types. There are also some approaches that use *Borda count* for ensemble methods (see for example the work of [19]). Moreover, voting rules have been applied to the specific case of Bagging [35, 36]. However, in [35], the authors combine only classifiers from the same family (i.e., Naive Bayes classifier) without mixing them.

A different perspective comes from the work of [15] and further improvements [11, 13, 55] where the basic assumption is that there always exists a correct ranking of the alternatives, but this cannot be observed directly. Voters derive their preferences over the alternatives from this ranking (perturbing it with noise). Scoring voting rules are proved to be maximum likelihood estimators (MLE). Under this approach, one computes the likelihood of the given preference profile for each possible state of the world, that is, the true ranking of the alternatives and the best ranking of the alternatives are then the ones that have the highest likelihood of producing the given profile. This model aligns very well with our proposal and justifies the use of voting rules in the aggregation of classifiers’ prediction. Moreover, MLEs give also a justification to the performance of ensembles where voting rules are used.

3 VORACE

The main idea of VORACE (VOTing with RANdom ClassifiERs) is that, given a sample, the output of each classifier can be seen as a ranking over the available classes, where the ranking order is given by the classifier's expected probability that the sample belongs to a class. Then a voting rule is used to aggregate these orders and declare a class as the "winner". VORACE generates a profile of n classifiers (where n is an input parameter) by randomly choosing the type of each classifier among a set of predefined ones. For instance, the classifier type can be drawn between a decision tree or a neural network. For each classifier, some of its hyper-parameters values are chosen at random, where the choice of which hyper-parameters and which values are randomly chosen depends on the type of the classifier. When all classifiers are generated, they are trained using the same set of training samples. For each classifier, the output is a vector with as many elements as the classes, where the i -th element represents the probability that the classifier assigns the input sample to the i -th class. Output values are ordered from the highest to the smallest one, and the output of each classifier is interpreted as a ranking over the classes, where the class with the highest value is the first in the ranking, then we have the class that has the second highest value in the output of the classifier, and so on. These rankings are then aggregated using a voting rule. The winner of the election is the class with the higher score. This corresponds to the prediction of VORACE. Ties can occur when more than one class gets the same score from the voting rule. In these cases, the winner is elected using a tie-breaking rule, which chooses the candidate that is most preferred by the classifier with the highest validation accuracy in the profile.

Example 1 Let us consider a profile composed by the output vectors of three classifiers, say y_1 , y_2 and y_3 , over four candidates (classes) c_1 , c_2 , c_3 and c_4 : $y_1 = [0.4, 0.2, 0.1, 0.3]$, $y_2 = [0.1, 0.3, 0.2, 0.4]$, and $y_3 = [0.4, 0.2, 0.1, 0.3]$. For instance, y_1 represents the prediction of the first classifier, which could predict that the input sample belongs to the first class with probability 0.4, to the second class with probability 0.2, to the third class with probability 0.1 and to the fourth class with probability 0.3. From the previous predictions we can derive the correspondent ranked orders x_1 , x_2 and x_3 . For instance, from prediction y_1 we can see that the first classifier prefers c_1 , then c_4 , then c_2 and then c_3 is the less preferred class for the input sample. Thus we have: $x_1 = [c_1, c_4, c_2, c_3]$, $x_2 = [c_4, c_2, c_3, c_1]$ and $x_3 = [c_1, c_4, c_2, c_3]$. Using Borda, class c_1 gets 6 points, c_2 gets 4 points, c_3 gets 1 point and c_4 gets 7 points. Therefore, c_4 is the winner, i.e. VORACE outputs c_4 as the predicted class. On the other hand, if we used Plurality, the winning class would be c_1 , since it is preferred by 2 out of 3 voters.

Notice that this method does not need any knowledge of architecture, type, or parameters, of the individual classifiers.²

² Code available at <https://github.com/aloreggia/vorace/>.

Table 1 Description of the datasets

Dataset	#Examples	#Categorical	#Numerical	Missing	#Classes
Anneal	898	32	6	yes	6
Autos	205	10	15	yes	7
Balance-s	625	0	4	no	3
Breast-cancer	286	9	0	yes	2
Breast-w	699	0	9	yes	2
Cars	1728	6	0	no	4
Credit-a	690	9	6	yes	2
Colic	368	15	7	yes	2
Dermatology	366	33	1	yes	6
Glass	214	0	9	no	5
Haberman	306	0	3	no	2
Heart-statlog	270	0	13	no	2
Hepatitis	155	13	6	yes	2
Ionosphere	351	34	0	no	2
Iris	150	0	4	no	3
Kr-vs-kp	3196	0	36	no	2
Letter	20,000	0	16	no	26
Lymphogra	148	15	3	no	4
Monks-3	122	6	0	no	2
Spambase	4,601	0	57	no	2
Vowel	990	3	10	no	11
Wine	178	0	13	no	3
Zoo	101	16	1	no	7

4 Experimental results

We considered 23 datasets from the UCI [43] repository. Table 1 gives a brief description of these datasets in terms of number of examples, number of features (where some features are categorical and others are numerical), whether there are missing values for some features, and number of classes. To generate the individual classifiers, we use three classification algorithms: Decision Trees (DT), Neural Networks (NN), and Support Vector Machines (SVM).

Neural networks are generated by choosing 2, 3 or 4 hidden layers with equal probability. For each hidden layer, the number of nodes is sampled geometrically in the range $[A, B]$, which means computing $\lfloor e^x \rfloor$ where x is drawn uniformly in the interval $[\log(A), \log(B)]$ [6]. We choose $A = 16$ and $B = 128$. The activation function is chosen with equal probability between the rectifier function $f(x) = \max(0, x)$ and the hyperbolic tangent function. The maximum number of epochs to train each neural network is set to 100. An early stopping callback is used to prevent the training phase to continue even when the accuracy is not improving and we set the patience parameter to $p = 5$. Batch size value is adjusted to respect the size of the dataset: given a training set T with size l , the batch size is set to $b = 2^{\lfloor \log_2(x) \rfloor}$ where $x = \frac{l}{100}$.

Decision trees are generated by choosing between the *entropy* and *gini* criteria with equal probability, and with a maximal depth uniformly sampled in [5, 25].

SVMs are generated by choosing randomly between the *rbf* and *poly* kernels. For both types, the *C* factor is drawn geometrically in $[2^{-5}, 2^5]$. If the type of the kernel is poly, the coefficient is sampled at random in [3, 5]. For rbf kernel, the gamma parameter is set to auto.

We used the average F1-score of a classifier ensemble as the evaluation metric, for all 23 different data sets, since the F1-score is a better measure to use if we need to seek a balance between Precision and Recall. For each dataset, we train and test the ensemble method with a 10-fold cross validation process. Additionally, for each dataset, experiments are performed 10 times, leading to a total of 100 runs for each method over each dataset. This is done to ensure greater stability. The voting rules considered in the experiments are Plurality, Borda, Copeland and Kemeny.

In order to compute the Kemeny consensus, we leverage the implementation of the Kemeny method for rank aggregation of incomplete rankings with ties that is available with the Python package named *corankco*³. The package provides several methods for computing a Kemeny consensus. Finding a Kemeny consensus is computationally hard, especially when the number of candidates grows. In order to ensure the feasibility of the experiments, we compute a Kemeny consensus using the exact algorithm with ILP Cplex when the number of classes $|C| \leq 5$, otherwise we employed the consensus computation with a heuristic (see package documentation for further details). We compare the performance of VORACE to 1) the average performance of a profile of individual classifiers, 2) the performance of the best classifier in the profile, 3) two state-of-the-art methods (Random Forest and XGBoost), and 4) the *Sum* method (also called *weighted averaging*). The *Sum* method computes $x_j^{\text{Sum}} = \sum_i^n x_{j,i}$ for each individual classifier *i* and for each class *j*, where $x_{j,i}$ is the probability that the sample belongs to class *j* predicted by classifier *i*. The winner is the one with the maximum value in the sum vector: $\arg \max_j x_j^{\text{Sum}}$. We did not compare VORACE to a more sophisticated version of Sum, such as *conditional averaging*, since they are not applicable in our case, requiring additional knowledge of the domain which is out of the scope of our work. Both Random Forest and XGBoost classifiers are generated with the same number of trees as the number of classifiers in the profile, all the remaining parameters are generated using default values. We did not compare to *stacking* because it would require to manually identify the correct structure of the sequence of classifiers in order to obtain competitive results. An optimal structure (i.e., a definition of a second level meta-classifier) can be defined by an expert in the domain at hand [8], and this is out of the scope of our work.

To study the accuracy of our method, we performed three kinds of experiments: 1) varying the number of individual classifiers in the profile and averaging the performance over all datasets, 2) fixing the number of individual classifiers and analyzing the performance on each dataset and 3) considering the introduction of more complex classifiers as base classifiers for VORACE. Since the first experiment shows that the best accuracy of the ensemble occurs when $n = 50$, we use only this size for the second and third experiments.

³ The package is available at <https://pypi.org/project/corankco/>.

Table 2 Average F1-scores (and standard deviation), varying the number of voters, averaged over all datasets

# Voters	Avg. profile	Borda	Plurality	Copeland	Kemeny	Sum	Best C.
5	0.8599 (0.1021)	0.8864 (0.1043)	0.8885 (0.1052)	0.8885 (0.1051)	0.8886 (0.1050)	0.8864 (0.1116)	0.8720 (0.1199)
7	0.8652 (0.0990)	0.8942 (0.0995)	0.8966 (0.1005)	0.8965 (0.1007)	0.8966 (0.1007)	0.8942 (0.1052)	0.8689 (0.1168)
10	0.8626 (0.0988)	0.8990 (0.0979)	0.9007 (0.0998)	0.9004 (0.1001)	0.9008 (0.1007)	0.8985 (0.1050)	0.8667 (0.1196)
20	0.8615 (0.0965)	0.9015 (0.0968)	0.9043 (0.0977)	0.9036 (0.0981)	0.9033 (0.0987)	0.8992 (0.1065)	0.8655 (0.1203)
40	0.8630 (0.0960)	0.9044 (0.0958)	0.9066 (0.0967)	0.9060 (0.0968)	0.9058 (0.0969)	0.9006 (0.1050)	0.8651 (0.1183)
50	0.8633 (0.0957)	0.9044 (0.0962)	0.9068 (0.0970)	0.9060 (0.0970)	0.9062 (0.0972)	0.8995 (0.1076)	0.8655 (0.1204)
Avg	0.8626 (0.0981)	0.8983 (0.0987)	0.9006 (0.0998)	0.9002 (0.0998)	0.9002 (0.1001)	0.8964 (0.1070)	0.8673 (0.1192)

4.1 Experiment 1: varying the number of voters in the ensemble

The aim of the first experiment is twofold: on one hand, we want to show that increasing the number of classifiers in the profile leads to an improvement of the performance. On the other hand, we want to show the effect of the aggregation on performance, compared with the best classifier in the profile and with the average classifier's performance. To do that, we first evaluate the overall average accuracy of VORACE varying the number n of individual classifiers in the profile. Table 2 presents the performance of each ensemble for different numbers of classifiers, specifically $n \in \{5, 7, 10, 20, 40, 50\}$. Plurality, Copeland, and Kemeny voting rules have their best accuracy for VORACE when $n = 50$. We set the system to stop the experiment after the time limit of one week, this is why we stop when $n = 50$. We are planning to run experiments with larger time limits in order to understand whether the system shows that the effect of the profile's size diminishes at some point. In Table 2, we report the F1-score and the standard deviation of VORACE with the considered voting rules. The last line of the table presents the average F1-score for each voting rule. The dataset "letter" was not considered in this test.

Increasing the number of classifiers in the ensemble, all the considered voting rules show an increase of the performance, specifically the higher the number of the classifiers the higher the F1-score of VORACE.

However, in Table 2 we can observe that the performance is slightly incremental when we increase the number of classifiers. This is due to the fact that in this particular experiment the accuracy of every single classifier is usually very high (i.e., $p \geq 0.8$), thus the ensemble has a reduced contribution to the aggregated result. In general this is not the case, especially when we have to deal with "harder" datasets where the accuracy p of single classifiers is lower. In Sect. 5, we will explore this case and we will see that the number of classifiers has a greater impact on the accuracy of the ensemble when the accuracy of the classifiers on average is low (e.g., $p \leq 0.6$).

Moreover, it is worth noting that the computational cost of the ensemble (both training and testing) increases linearly with the number of classifiers in the profile. Thus, it is convenient to consider more classifiers, especially when the accuracy of the single classifiers is poor. Thus, overall, the increase in the number of classifiers has a positive effect on the performance of VORACE, as expected given the theoretical analysis in Sect. 5⁴.

For each voting rule, we also compared VORACE to the average performance of the individual classifiers and the best classifier in the profile, to understand if VORACE is better than the best classifier, or if it is just better than the average classifiers' accuracy (around 0.86). In Table 2 we can see that VORACE always behaves better than both the best classifier and the profile's average. Moreover, it is interesting to observe that Plurality performs better on average than more complex voting rules like Borda and Copeland.

4.2 Experiment 2: comparing with existing methods

For the second experiment, we set $n = 50$ and we compare VORACE (using Majority, Borda, Plurality, Copeland, and Kemeny) with Sum, Random Forest (RF), and XGBoost in each dataset. Table 3 reports the performances of VORACE on binary datasets where all the voting rules collapse to Majority voting. VORACE performances are very close to

⁴ However, the experiments do not satisfy the independence assumption of the theoretical study

Table 3 Performances on binary datasets: average F1-scores (and standard deviation)

Dataset	Majority	Sum	RF	XGBoost
Breast-cancer	0.7356 (0.0947)	0.7151 (0.0983)	0.7134 (0.0397)	0.7000 (0.0572)
Breast-w	0.9645 (0.0133)	0.9610 (0.0168)	0.9714 (0.0143)	0.9613 (0.0113)
Colic	0.8587 (0.0367)	0.8573 (0.0514)	0.8507 (0.0486)	0.8750 (0.0534)
Credit-a	0.8590 (0.0613)	0.8478 (0.0635)	0.8710 (0.0483)	0.8565 (0.0763)
Haberman	0.7337 (0.0551)	0.6994 (0.0765)	0.7353 (0.0473)	0.7158 (0.0518)
Heart-statlog	0.8070 (0.0699)	0.7885 (0.0797)	0.8259 (0.0621)	0.8222 (0.0679)
Hepatitis	0.8385 (0.0903)	0.8377 (0.0955)	0.8446 (0.0610)	0.8242 (0.0902)
Ionosphere	0.9435 (0.0348)	0.9366 (0.0344)	0.9344 (0.0385)	0.9260 (0.0427)
Kr-vs-kp	0.9958 (0.0044)	0.9960 (0.0044)	0.9430 (0.0139)	0.9562 (0.0174)
Monks-3	0.9182 (0.0712)	0.9115 (0.0748)	0.9333 (0.0624)	0.9333 (0.0624)
Spambase	0.9416 (0.0105)	0.8801 (0.1286)	0.9100 (0.0137)	0.9294 (0.0112)
Average	0.8724 (0.0493)	0.8574 (0.0658)	0.8666 (0.0409)	0.8636 (0.0493)

Best performance in bold. On binary datasets, all the voting rules behave as majority voting rule

the state-of-the-art. We try to use Kemeny on the dataset “letter” but it exceeds the time limit of one week and thus it was not possible to compute the average. In order to make the average values comparable (last row of Table 4), performances on the dataset “letter” were not considered in the computation of the average values for the other methods. Table 4 reports the performances on datasets that have multiple classes: when the number of classes increases VORACE is still stable and behaves very similarly to the state-of-the-art. The similarity among the performances is promising for the system. Indeed, RandomForest and XGBoost reach better performances on some datasets and they can be improved on over by optimizing their hyperparameters. But, this experiment shows that it is possible to reach very similar performances using a very simple method as VORACE is. This means that usage of VORACE does not require any optimization of hyperparameters whether it is done manually or automatically. The importance of this property is corroborated by a recent line of work by [52] that suggests how industry and academia should focus their efforts on developing tools that reduce or avoid hyperparameters’ optimization, resulting in simpler methods that are also more sustainable in terms of energy and time consumption.

Moreover, Plurality is both more time and space efficient since it needs a smaller amount of information: for each classifier it just needs the most preferred candidate instead of the whole ranking, contrarily to other methods such as Sum. We also performed two additional variants of these experiments, one with a weighted version of the voting rules (where the weights are the classifiers’ validation accuracy), and the other one by training each classifier on different portions of the data in order to increase the independence between them. In both experiments, the results are very similar to the ones reported here.

4.3 Experiment 3: introducing complex classifiers in the profile

The goal of the third experiment is to understand whether using complex classifiers in the profile (such as using an ensemble of ensembles) would produce better final performances. For this purpose, we compared VORACE with standard base classifiers (described in Sect. 3) with three different versions of VORACE with complex base classifiers: 1)

Table 4 Performances on multiclass datasets: average F1-scores (and standard deviation)

Dataset	Borda	Plurality	Copeland	Kemeny	Sum	RF	XGBoost
Anneal	0.9917 (0.0138)	0.9876 (0.0200)	0.9876 (0.0200)	0.9880 (0.0194)	0.9894 (0.0174)	0.8471 (0.0122)	0.9912 (0.0080)
Autos	0.8021 (0.0669)	0.7848 (0.0794)	0.7803 (0.0768)	0.7832 (0.0771)	0.8095 (0.0749)	0.6890 (0.0743)	0.8298 (0.0744)
Balance	0.9016 (0.0366)	0.9208 (0.0311)	0.9069 (0.0292)	0.9082 (0.0297)	0.8911 (0.0376)	0.8561 (0.0540)	0.8578 (0.0441)
Cars	0.9916 (0.0079)	0.9932 (0.0054)	0.9931 (0.0054)	0.9934 (0.0053)	0.9931 (0.0048)	0.7928 (0.0300)	0.8935 (0.0266)
Dermatology	0.9819 (0.0192)	0.9769 (0.0206)	0.9769 (0.0206)	0.9766 (0.0209)	0.9783 (0.0196)	0.9699 (0.0256)	0.9755 (0.0189)
Glass	0.9708 (0.0364)	0.9602 (0.0319)	0.9607 (0.0291)	0.9611 (0.0287)	0.9742 (0.0268)	0.9535 (0.0295)	0.9719 (0.0313)
Iris	0.9473 (0.0576)	0.9473 (0.0576)	0.9473 (0.0576)	0.9480 (0.0570)	0.9527 (0.0519)	0.9533 (0.0521)	0.9600 (0.0442)
Letter	0.9311 (0.01)	0.9590 (0.01)	0.9545 (0.01)	–	0.9627 (0.01)	0.6044 (0.01)	0.8832 (0.01)
Lymphography	0.8461 (0.0983)	0.8630 (0.0851)	0.8624 (0.0843)	0.8604 (0.0875)	0.8529 (0.0925)	0.8586 (0.0691)	0.8519 (0.0490)
Vowel	0.9476 (0.0232)	0.9862 (0.0110)	0.9860 (0.0116)	0.9860 (0.0114)	0.9862 (0.0119)	0.7333 (0.0323)	0.8323 (0.0333)
Wine	0.9656 (0.0537)	0.9789 (0.0331)	0.9783 (0.0342)	0.9783 (0.0342)	0.9806 (0.0380)	0.9889 (0.0222)	0.9611 (0.0558)
Zoo	0.9550 (0.0497)	0.9550 (0.0517)	0.9560 (0.0496)	0.9590 (0.0492)	0.9500 (0.0640)	0.9500 (0.0500)	0.9700 (0.0640)
Average	0.9365 (0.0421)	0.9413 (0.0388)	0.9396 (0.0380)	0.9402 (0.0382)	0.9416 (0.0399)	0.8720 (0.0410)	0.9177 (0.0409)

Best performance in bold

Table 5 Average F1-scores (and standard deviation)

dataset	VORACE with RF & XGBoost	VORACE with only RF	VORACE with only XGBoost
Anneal	0.9937 (0.01)	0.9921 (0.01)	0.9893 (0.01)
Autos	0.8095 (0.09)	0.7969 (0.10)	0.7916 (0.08)
Balance	0.8998 (0.02)	0.8456 (0.03)	0.8040 (0.04)
Breast-cancer*	0.7573 (0.04)	0.7485 (0.06)	0.7394 (0.06)
Breast-w*	0.9654 (0.02)	0.9744 (0.02)	0.9605 (0.03)
Cars	0.9887 (0.01)	0.9547 (0.01)	0.9044 (0.05)
Colic*	0.8668 (0.04)	0.8766 (0.04)	0.8638 (0.04)
Credit-a*	0.8737 (0.03)	0.8691 (0.03)	0.8712 (0.03)
Dermatology	0.9749 (0.02)	0.9765 (0.02)	0.9805 (0.02)
Glass	0.9761 (0.03)	0.9740 (0.04)	0.9770 (0.03)
Haberman*	0.7338 (0.04)	0.7168 (0.04)	0.7286 (0.02)
Heart-statlog*	0.8315 (0.09)	0.8352 (0.09)	0.8248 (0.08)
Hepatitis*	0.8215 (0.07)	0.8091 (0.05)	0.8105 (0.08)
Ionosphere*	0.9349 (0.04)	0.9272 (0.05)	0.9347 (0.04)
Iris	0.9627 (0.05)	0.9593 (0.04)	0.9593 (0.05)
Kr-vs-kp*	0.9953 (0.00)	0.9869 (0.01)	0.9892 (0.01)
Letter	0.9632 (0.01)	0.9622 (0.01)	0.9265 (0.01)
Lymphography	0.8700 (0.10)	0.8306 (0.15)	0.8412 (0.14)
Monks-3*	0.9156 (0.07)	0.9340 (0.06)	0.9037 (0.07)
Spambase*	0.9437 (0.01)	0.9439 (0.01)	0.9337 (0.01)
Vowel	0.9834 (0.01)	0.9691 (0.02)	0.9086 (0.03)
Wine	0.9851 (0.03)	0.9764 (0.04)	0.9796 (0.04)
Zoo	0.9535 (0.05)	0.9589 (0.05)	0.9231 (0.06)
Average	0.9130 (0.04)	0.9051 (0.04)	0.8933 (0.04)

* Denotes binary datasets

VORACE with only Random Forest 2) VORACE with only XGBoost and 3) VORACE with Random Forest, XGBoost and standard base classifiers (DT, SVM, NN).

For simplicity, we used the Plurality voting rule, since it is the most efficient method and it is one of the voting rules that gives better results. We fixed the number of voters in the profiles to 50 and we selected the parameters for the simple classifiers for VORACE as described at the beginning of Sect. 4. For Random Forest, parameters were drawn uniformly among the following values⁵: *bootstrap* between *True* and *False*, *max_depth* in [10, 20, ..., 100, *None*], *max_features* between [*auto*, *sqrt*], *min_samples_leaf* in [1, 2, 4], *min_samples_split* in [2, 5, 10], and *n_estimators* in [10, 20, 50, 100, 200]. For XGBoost instead the parameters were drawn uniformly among the following values: *max_depth* in the range [3, 25], *n_estimators* equals the number of classifiers, *subsample* in [0, 1], and *colsample_bytree* in [0, 1]. The results of the comparison between the different versions of VORACE are provided in Table 5. We can observe that the performance of VORACE (column “Majority” of Table 3 and column “Plurality” of Table 4) is not significantly improved

⁵ Parameters’ names and values refer to the Python’s modules: `RandomForestClassifier` in `sklearn.ensemble` and `xgb` in `xgboost`.

by using more complex classifiers as a base for the profile. It is interesting to notice the effect of VORACE on the aggregation of RF with respect to a single RF. Comparing the results in Table 3 and 4 (RF column) with results in Table 5 (VORACE with only RF column), one can notice that RF is positively affected by the aggregation on many datasets (on all the datasets the improvement is on average 5%), especially on those with multiple classes. Moreover, the improvement is significant in many of them: i.e. on “letter” dataset we have an improvement of more than 35%. This effect can be explained by the random aggregation of trees used by the RF algorithm, where the goal is to reduce the variance of the single classifier. In this sense, a principled aggregation of different RF models (as the one in VORACE) is a correct way to boost the final performance: distinct RF models act differently over separate parts of the domain, providing VORACE with a good set of weak classifiers—see Theorem 3.

We saw in this section that this more complex version of VORACE does not provide any significant advantage, in terms of performance, compared with the standard one. To conclude, we thus suggest using VORACE in its standard version without adding complexity to the base classifiers.

In other experiments, we also see that the probability of choosing the correct class decreases if the number of classes increases. This means that the task becomes more difficult with a larger number of classes.

5 Theoretical analysis: independent classifiers with the same accuracy

In this section we theoretically analyze the probability to predict the correct label/class of our ensemble method.

Initially, we consider a simple scenario with m classes (the candidates) and a profile of n independent classifiers (the voters), where each classifier has the same probability p of correctly classifying a given instance. The independence assumption hardly fully holds in practice, but it is a natural simplification (commonly adopted in literature) used for the sake of analysis.

We assume that the system uses the Plurality voting rule. This is justified by the fact that Plurality provides better results in our experimental analysis (see Sect. 4) and thus it is the one we suggest to use with VORACE. Moreover, Plurality has also the advantage to require very little information from the individual classifiers and also being computationally efficient.

We are interested in computing the probability that VORACE chooses the correct class. This probability corresponds to the accuracy of VORACE when considering the single classifiers as black boxes, i.e., knowing only their accuracy without any other information. The result presented in the following theorem is especially powerful because it shows a closed formula that only requires for the values of p , m , and n to be known.

Theorem 1 *The probability of electing the correct class c^* , among m classes, with a profile of n classifiers, each one with accuracy $p \in [0, 1]$, using Plurality is given by:*

$$\mathcal{T}(p) = \frac{1}{K} (1-p)^n \sum_{i=\lceil \frac{n}{m} \rceil}^n \varphi_i(n-i)! \binom{n}{i} \left(\frac{p}{1-p} \right)^i \quad (1)$$

where φ_i is defined as the coefficient of the monomial x^{n-i} in the expansion of the following generating function:

$$\mathcal{G}_i^m(x) = \left(\sum_{j=0}^{i-1} \frac{x^j}{j!} \right)^{m-1}$$

and K is a normalization constant defined as:

$$K = \sum_{j=0}^n \binom{n}{j} p^j (m-1)^{n-j} (1-p)^{n-j}.$$

Proof The formula can be rewritten as:

$$\mathcal{T}(p) = \frac{1}{K} \sum_{i=\lfloor \frac{n}{m} \rfloor}^n \binom{n}{i} p^i \varphi_i(n-i)! (1-p)^{n-i}$$

and corresponds to the sum of the probability of all the possible different profiles votes that elect c^* . We perform the sum varying i , an index that indicates the number of classifiers in the profile that vote for the correct label c^* . This number is between $\lfloor \frac{n}{m} \rfloor$ (since if $i < \lfloor \frac{n}{m} \rfloor$ that profile cannot elect c^*) and n where all the classifiers vote for c^* . The binomial factor expresses the number of possible positions, in the ordered profile of size n , of i classifiers that votes for c^* . This is multiplied by the probability of these classifiers to vote c^* , that is p^i . The factors $\varphi_i(n-i)!$ correspond the number of possible combinations of votes of the $n-i$ classifiers (on the other candidates different from c^*) that ensure the winning of c^* . This is computed as the number of possible combinations of $n-i$ objects extracted from a set $(m-1)$ objects, with a bounded number of repetitions (bounded by $i-1$ to ensure the winning of c^*). The formula to use for counting the number of combinations of D objects extracted from a set A objects, with a bounded number of repetitions B , is: $\varphi_i D!$. In our case $A = m-1$ is the number of objects, $B = i-1$ is the maximum number of repetitions and $D = n-i$ the positions to fill and φ_i is the coefficient of x^D in the expansion of the following generating function:

$$\left(\sum_{j=0}^B \frac{x^j}{j!} \right)^A \xrightarrow[B=i-1]{A=m-1} \left(\sum_{j=0}^{i-1} \frac{x^j}{j!} \right)^{m-1} = \mathcal{G}_i^m(x).$$

Finally, the factor $(1-p)^{n-i}$ is the probability that the remaining $n-i$ classifiers do not elect c^* . □

For the sake of comprehension, we give an example that describes the computation of the probability of electing the correct class c^* , as formalized in Theorem 1.

Example 2 Considering an ensemble with 3 classifiers (i.e., $n = 3$), each one with accuracy $p = 0.8$. The number of classes in the dataset is $m = 4$. The probability of choosing the correct class c^* is given by the formula in Theorem 1. Specifically:

$$\mathcal{T}(p) = (1 - 0.8)^3 \frac{1}{K} \sum_{i=1}^3 \varphi_i \cdot (3 - i)! \binom{3}{i} \left(\frac{0.8}{1 - 0.8} \right)^i$$

where $K = 1.728$. In order to compute the value of each φ_i , we have to compute the coefficient of x^{3-i} in the expansion of the generating function $\mathcal{G}_i^4(x)$.

For $i = 1$: We have $\mathcal{G}_1^4(x) = 1$, and we are interested in the coefficient of $x^{n-i} = x^2$, thus $\varphi_1 = 0$.

For $i = 2$: We have $\mathcal{G}_2^4(x) = 1 + 3x + 3x^2 + x^3$, and we are interested in the coefficient of $x^{n-i} = x^1$, thus $\varphi_2 = 3$.

For $i = 3$: We have $\mathcal{G}_3^4(x) = 1 + 3x + \frac{9}{2}x^2 + 4x^3 + \frac{9}{4}x^4 + \frac{3}{4}x^5 + \frac{1}{8}x^6$, and we are interested in the coefficient of $x^{n-i} = x^0$, thus $\varphi_3 = 1$.

We can now compute the probability $\mathcal{T}(p)$:

$$\begin{aligned} \mathcal{T}(p) &= \frac{0.008}{1.728} \cdot (\varphi_1 \cdot (2)! \binom{3}{1} \cdot 4 + \varphi_2 \cdot (1)! \binom{3}{2} \cdot (4)^2 + \varphi_3 \cdot (0)! \binom{3}{3} \cdot (4)^3) \\ &= 0.963. \end{aligned}$$

The result says that VORACE with 3 classifiers (each one with accuracy $p = 0.8$) has a probability of 0.963 of choosing the correct class c^* .

It is worth noting that $\mathcal{T}(p) = 1$ when $p = 1$, meaning that, when all the classifiers in the ensemble always predict the right class, our ensemble method always outputs the correct class as well⁶. Moreover, $\mathcal{T}(p) = 0$ in the symmetric case in which $p = 0$, that is when all the classifiers always predict a wrong class.

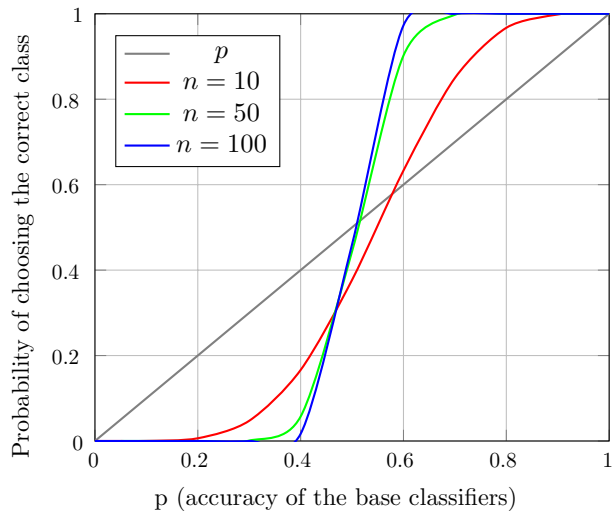
Note that the independence assumption considered above is in line with previous studies (e.g., the same assumption is made in [10, 55]) and it is a necessary simplification to obtain a closed formula to compute $\mathcal{T}(p)$. Moreover, in a realistic scenario, p can be interpreted as representing the lower bound of the accuracy of the classifiers in the profile. It is easy to see that under this interpretation the value of $\mathcal{T}(p)$ as well represents a lower bound of the probability of electing the correct class c^* , given m available classes, and a profile of n classifiers.

Although this theoretical result holds in a restricted scenario and with a specific voting rule, as we already noticed in our experimental evaluation in Sect. 4, the probability of choosing the correct class is always greater than or equal to each individual classifiers' accuracy.

It is worth noting that the scenario considered above is similar to the one analyzed in the *Condorcet Jury Theorem* [10], which states that in a scenario with two candidates where each voter has probability $p > \frac{1}{2}$ to vote for the correct candidate, the probability that the correct candidate is chosen goes to 1 as the number of votes goes to infinity. Some restrictions imposed by this theorem are partially satisfied also in our scenario: some voters (classifiers) are independent on each other (those that belong to a different classifier's category),

⁶ Formula 1 is equal to 1 for $p = 1$ because all the terms of the sum are equal to zero except the last term for $i = n$ ($K = 1$ and $\varphi_i(0) = 1$ as well). This is equal to 1 because we have $(1 - p)^0 = 0^0$ and by convention $0^0 = 1$ when we are considering discrete exponents.

Fig. 1 Probability of choosing the correct class c^* varying the size of the profile $n \in \{10, 50, 100\}$ and keeping m constant to 2, where each classifier has the same probability p of classifying a given instance correctly



since we generate them randomly. However, Theorem 1 does not immediately follow from this result. Indeed, it represents a generalization because some of the Condorcet restrictions do not hold in our case, specifically: 1) 2-class classification task does not hold, since VORACE can be used also with more than 2 classes; 2) classifiers are generated randomly, thus we cannot ensure that the accuracy $p > \frac{1}{2}$, especially with more than two classes. This work has been reinterpreted first by [55] and successively extended by [44] and [51], considering the cases in which the agents/voters have different p_i . However, the focus of these works is fundamentally different from ours, since their goal is to find the optimal decision rule that maximizes the probability that a profile elects the correct class.

Given the different conditions of our setting, we cannot apply the Condorcet Jury Theorem, or the works cited above, as such. However, in Sect. 6 we will formally see that considering $m = 2$, our formulation enforces the results stated by Condorcet Jury Theorem.

Moreover, our work is in line with the analysis regarding maximum likelihood estimators (MLEs) for r -noise models [11, 50]. An r -noise model is a noise model for ranking over a set of candidates, i.e., a family of probability distributions in the form $P(\cdot|u)$, where u is the correct preference. This means that an r -noise model describes a voting process where there is a ground truth about the collective decision, although the voters do not know it. In this setting, a MLE is a preference aggregation function f that maximizes the product of the probabilities $P(v_i|u)$, $i = 1, \dots, n$ for a given voters' profile $R = (v_1, \dots, v_n)$. Finding a suitable f corresponds to our goal.

MLEs for r -noise models have been studied in details by [11] assuming the noise is independent across votes. This corresponds to our preliminary assumption of the independence of the base classifiers. The first result in [11] states that given a voting rule, there always exists a r -noise model such that the voting rule can be interpreted as a MLE (see Theorem 1 in [11]). In fact, given an appropriate r -noise model, any scoring rule is a maximum likelihood estimator for winner under i.i.d. votes. Thus, for a given input sample, we can interpret the classifiers rankings as a permutation of the true ranking over the classes and the voting rule (like Plurality and Borda) used to aggregate these rankings as a MLE for an r -noise model on the original classification of the examples. However, to the best of our knowledge, providing a closed formulation (i.e., considering only the main problem's

parameters p , m and n , and without having any information on the original true ranking or the noise model) to compute the probability of electing the winner (as provided in our Theorem 1) for a given profile using Plurality is a novel and valuable contribution (see discussion on the attempts existing in the literature to define the formula in Sect. 2.3). We remind the reader that in our learning scenario the formula in Theorem 1 is particularly useful because it computes a lower bound on the accuracy of VORACE (that is, the probability that VORACE selects the correct class) when knowing only the accuracy of the base classifiers, considering them as black boxes.

More precisely, we analyze the relationship between the probability of electing the winner (i.e., Formula 1) and the accuracy of each individual classifier p . Figure 1 shows the probability of choosing the correct class varying the size of the profile $n \in \{10, 50, 100\}$ and keeping $m = 2$.⁷ We see that, by augmenting the size of the profile n , the probability that the ensemble chooses the right class grows as well. However, the benefit is just incremental when base classifiers have high accuracy. We can see that when p is high we reach a plateau where $\mathcal{T}(p)$ is very close to 1 regardless of the number of classifiers in the profile. In a realistic scenario, having a high baseline accuracy in the profile is not to be expected, especially when we consider “hard” datasets and randomly generated classifiers. In these cases (when the accuracy of the base classifiers in average is low), the impact of the number of classifiers is more evident (for example when $p = 0.6$).

Thus, if $p > 0.5$ and n tends to infinity, then it is beneficial to use a profile of classifiers. This is in line with the result of the *Condorcet Jury Theorem*.

6 Theoretical analysis: comparison with Condorcet Jury Theorem

In this section we prove how, for $m = 2$, Formula 1 enforces the results stated in the Condorcet Jury Theorem [10] (see Sect. 5 for the Condorcet Jury Theorem statement). Notice, as for Theorem 1, the adopted assumptions likely do not fully hold in practice, but are natural simplifications used for the sake of analysis. Specifically, we need to prove the following theorem.

Theorem 2 *The probability of electing the correct class c^* , among 2 classes, with a profile of an infinite number of classifiers, each one with accuracy $p \in [0, 1]$, using Plurality, is given by:*

$$\lim_{n \rightarrow \infty} \mathcal{T}(p) = \begin{cases} 0 & p < 0.5 \\ 0.5 & p = 0.5 \\ 1 & p > 0.5 \end{cases} \quad (2)$$

In Fig. 2 we can see a visualization of the function $\mathcal{T}(p)$ when $n \rightarrow \infty$, as described in Theorem 2. In what follows we will prove this by showing that the function $\mathcal{T}(p)$ is monotonic increasing and when $n \rightarrow \infty$ is equal to 0.

⁷ Figure 1 has been created by grid sampling the values of $p \in [0, 1]$ with step 0.05 and by performing an exact computation of the value of $\mathcal{T}(p)$ for each specific value of p in the sampling set with $n \in \{10, 50, 100\}$ and $m = 2$. We then connected these values with the smoothing algorithm of *TikZ* package.

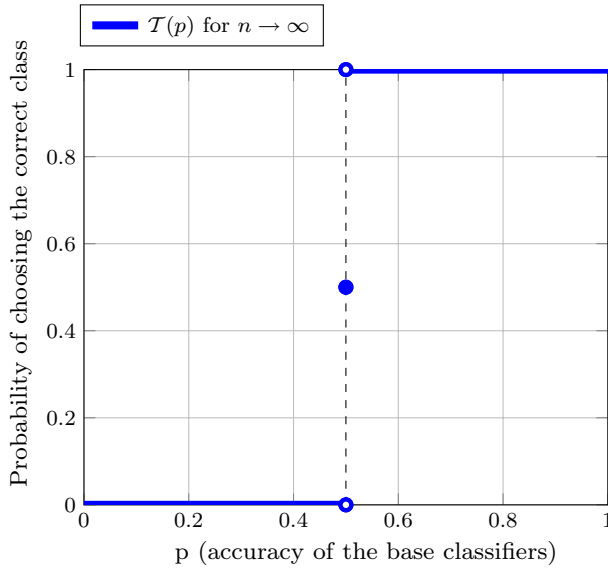


Fig. 2 The probability of electing the correct class c^* , among 2 classes, with a profile of an infinite number of classifiers ($n \rightarrow \infty$), each one with accuracy $p \in [0, 1]$, using Plurality

Firstly, we find an alternative, more compact, formulation for $\mathcal{T}(p)$ in the case of binary datasets (only two alternatives/candidates, i.e., $m = 2$) in the following Lemma.

Lemma 1 *The probability of electing the correct class c^* , among 2 classes, with a profile of n classifiers, each one with accuracy $p \in [0, 1]$, using Plurality is given by:*

$$\mathcal{T}(p) = \sum_{i=\lceil \frac{n}{2} \rceil}^n \binom{n}{i} p^i (1-p)^{n-i}. \tag{3}$$

Proof It is possible to note how for $m = 2$, the values of φ_i is $\frac{1}{(n-i)!}$.

This because:

$$\begin{aligned} \mathcal{G}_i^2(x) &= \left(\sum_{j=0}^{i-1} \frac{x^j}{j!} \right) \\ &= 1 + x + \frac{1}{2}x^2 + \dots + \frac{1}{(n-i)!}x^{n-i} + \dots + \frac{1}{(i-1)!}x^{i-1}. \end{aligned}$$

Consequently, with further algebraic simplifications, we have the following:

$$\begin{aligned}
 \mathcal{T}(p) &= \frac{1}{K}(1-p)^n \sum_{i=\lceil \frac{n}{2} \rceil}^n \varphi_i(n-i)! \binom{n}{i} \left(\frac{p}{1-p}\right)^i \\
 &= \frac{1}{K}(1-p)^n \sum_{i=\lceil \frac{n}{2} \rceil}^n \frac{(n-i)!}{(n-i)!} \binom{n}{i} \left(\frac{p}{1-p}\right)^i \\
 &= \frac{1}{K}(1-p)^n \sum_{i=\lceil \frac{n}{2} \rceil}^n \binom{n}{i} \left(\frac{p}{1-p}\right)^i \\
 &= \frac{(1-p)^n \sum_{i=\lceil \frac{n}{2} \rceil}^n \binom{n}{i} \left(\frac{p}{1-p}\right)^i}{\sum_{j=0}^n \binom{n}{j} p^j (1-p)^{n-j}} \\
 &= \frac{(1-p)^n \sum_{i=\lceil \frac{n}{2} \rceil}^n \binom{n}{i} \left(\frac{p}{1-p}\right)^i}{(1-p)^n \sum_{j=0}^n \binom{n}{j} \left(\frac{p}{1-p}\right)^j} \\
 \Rightarrow \mathcal{T}(p) &= \frac{\sum_{i=\lceil \frac{n}{2} \rceil}^n \binom{n}{i} \left(\frac{p}{1-p}\right)^i}{\sum_{i=0}^n \binom{n}{i} \left(\frac{p}{1-p}\right)^i}.
 \end{aligned}$$

Now, looking at the denominator, by definition of binomial coefficient, we can note that:

$$\sum_{i=0}^n \binom{n}{i} \left(\frac{p}{1-p}\right)^i = \left(1 + \frac{p}{1-p}\right)^n = (1-p)^{-n}.$$

Thus, we obtain:

$$\mathcal{T}(p) = \sum_{i=\lceil \frac{n}{2} \rceil}^n \binom{n}{i} p^i (1-p)^{n-i}.$$

□

We will now consider the two cases separately: (i) $p = 0.5$, and (ii) $p > 0.5$ or $p < 0.5$. For both cases we will prove the corresponding statement of Theorem 2.

6.1 Case: $p = 0.5$

We will now proceed to prove the second statement of Theorem 2.

Proof If $p = 0.5$ we have that:

$$\mathcal{T}(0.5) = \frac{\sum_{i=\lceil \frac{n}{2} \rceil}^n \binom{n}{i}}{2^n}.$$

We note that, if n is an odd number:

$$\sum_{i=\lceil \frac{n}{2} \rceil}^n \binom{n}{i} = \frac{\sum_{i=0}^n \binom{n}{i}}{2} = 2^{n-1},$$

while if n is even:

$$\sum_{i=\lceil \frac{n}{2} \rceil}^n \binom{n}{i} = \frac{\sum_{i=0}^n \binom{n}{i}}{2} + \frac{1}{2} \binom{n}{\frac{n}{2}}.$$

Thus, we have the two following cases, depending on n :

$$\mathcal{T}(0.5) = \frac{2^{n-1}}{2^n} = 0.5, \quad \text{if } n \text{ is odd;} \tag{4}$$

$$\mathcal{T}(0.5) = \frac{2^{n-1} + \frac{1}{2} \binom{n}{\frac{n}{2}}}{2^n} = 0.5 + \frac{\frac{1}{2} \binom{n}{\frac{n}{2}}}{2^n}, \quad \text{if } n \text{ is even.} \tag{5}$$

We can see that, when n is odd, the following term becomes 0 if n tend to infinity:

$$\lim_{n \rightarrow \infty} \frac{\frac{1}{2} \binom{n}{\frac{n}{2}}}{2^n} = \lim_{n \rightarrow \infty} \frac{\binom{n}{\frac{n}{2}}}{2^{n+1}} = 0.$$

This limit is an indeterminate form $\frac{\infty}{\infty}$, that can be easily solved considering that $\binom{n}{\frac{n}{2}} < 2^n$.

Given this observation we can see that the denominator prevails making the limit going to 0. Thus, we proved that:

$$\lim_{n \rightarrow \infty} \mathcal{T}(0.5) = 0.5.$$

□

We note that if n is odd $\mathcal{T}(0.5) = 0.5$ also for small values of n , while if n is even, $\mathcal{T}(0.5)$ converges to 0.5 and it is equal to 0.5 only when $n \rightarrow \infty$.

6.2 Monotonicity and analysis of the derivative

In this section, we first show that $\mathcal{T}(p)$ (see Equation 3) is monotonic increasing by proving that its derivative is greater or equal to zero. Finally, we will see that, at the limit (for $n \rightarrow \infty$), the derivative is equal to zero for every $p \in [0, 1]$ excluding 0.5.

Lemma 2 *The function $\mathcal{T}(p)$, describing the probability of electing the correct class c^* , among 2 classes, with a profile of a n classifiers, each one with accuracy $p \in [0, 1]$, using Plurality is monotonic increasing.*

Proof We know from Equation 3 in Lemma 1 that

$$\mathcal{T}(p) = \sum_{i=\lceil \frac{n}{2} \rceil}^n \binom{n}{i} p^i (1-p)^{n-i}.$$

We want now to prove that $\mathcal{T}(p) \geq 0$.

$$\begin{aligned} \frac{\partial \mathcal{T}(p)}{\partial p} &= \frac{\partial \left(\sum_{i=\lceil \frac{n}{2} \rceil}^n \binom{n}{i} p^i (1-p)^{n-i} \right)}{\partial p} \\ &= \sum_{i=\lceil \frac{n}{2} \rceil}^n \binom{n}{i} \frac{\partial (p^i (1-p)^{n-i})}{\partial p} \\ &= \sum_{i=\lceil \frac{n}{2} \rceil}^n \binom{n}{i} \left(\frac{\partial (p^i)}{\partial p} (1-p)^{n-i} + p^i \frac{\partial ((1-p)^{n-i})}{\partial p} \right) \\ &= \sum_{i=\lceil \frac{n}{2} \rceil}^n \binom{n}{i} (i p^{i-1} (1-p)^{n-i} - p^i (n-i) (1-p)^{n-i-1}) \\ &= \sum_{i=\lceil \frac{n}{2} \rceil}^n \binom{n}{i} p^{i-1} (1-p)^{n-i-1} (i - pi - pn + pi) \\ &= \sum_{i=\lceil \frac{n}{2} \rceil}^n \binom{n}{i} p^{i-1} (1-p)^{n-i-1} (i - pn) \\ &= (1-p)^{n-1} \sum_{i=\lceil \frac{n}{2} \rceil}^n \binom{n}{i} p^{i-1} (1-p)^{-i} (i - pn) \\ &= (1-p)^{n-1} \sum_{i=\lceil \frac{n}{2} \rceil}^n \binom{n}{i} \left(\frac{p}{1-p} \right)^{i-1} \left(\frac{i}{p} - n \right) \\ &= (1-p)^{n-1} \left(\frac{1-p}{p} \right)^{\lceil \frac{n}{2} \rceil} \binom{n}{\lceil \frac{n}{2} \rceil} \left(\frac{p}{1-p} \right)^{\lceil \frac{n}{2} \rceil} \\ &= p^{\lceil \frac{n}{2} \rceil - 1} (1-p)^{n - \lceil \frac{n}{2} \rceil} \binom{n}{\lceil \frac{n}{2} \rceil} \left(\frac{p}{1-p} \right) \geq 0 \end{aligned}$$

It is easy to see that the last row of the sequence is greater or equal to zero since each of the terms of the product is greater or equal to zero. We proved that $\mathcal{T}(p)$ is monotonic increasing. □

Let's see now that at the limit (with $n \rightarrow \infty$) the derivative is equal to zero for every $p \in [0, 1]$ excluding $p = 0.5$.

Lemma 3 Given the function $\mathcal{T}(p)$ describing the probability of electing the correct class c^* , among 2 classes, with a profile of a n classifiers, each one with accuracy $p \in [0, 1]$, using Plurality, we have that:

$$\lim_{n \rightarrow \infty} \frac{\partial \mathcal{T}(p)}{\partial p} = 0$$

Proof Let's rewrite the function $\frac{\partial \mathcal{T}(p)}{\partial p}$ as follows:

$$\frac{\partial \mathcal{T}(p)}{\partial p} = p^{\lceil \frac{n}{2} \rceil - 1} (1 - p)^{n - \lceil \frac{n}{2} \rceil} \binom{n}{\lceil \frac{n}{2} \rceil}$$

We will treat separately the case in which n is an odd or even number:

$$\frac{\partial \mathcal{T}(p)}{\partial p} = \begin{cases} (p(1 - p))^{\lfloor \frac{n}{2} \rfloor} \binom{n}{\lfloor \frac{n}{2} \rfloor} & p \text{ odd} \\ \frac{(p(1 - p))^{\frac{n}{2}}}{p} \binom{n}{\frac{n}{2}} & p \text{ even} \end{cases}$$

Case 1: n is odd. This is an indeterminate form $0 \cdot \infty$, that can be solved considering that:

$$[p(1 - p)]^{\lfloor \frac{n}{2} \rfloor} \leq \frac{\partial \mathcal{T}(p)}{\partial p} \leq (2 + \frac{1}{n})^n (p(1 - p))^{\lfloor \frac{n}{2} \rfloor}$$

where the inequality on the right follows from:

$$1 < \binom{n}{\lfloor \frac{n}{2} \rfloor} < (2 + \frac{1}{n})^n.$$

Let's consider the function of the left inequality when $n \rightarrow \infty$.

Since $p(1 - p) < 1 \forall p \in [0, 1]$, we know that:

$$\lim_{n \rightarrow \infty} [p(1 - p)]^{\lfloor \frac{n}{2} \rfloor} = 0$$

This can be proved with the following observation:

$$p(1 - p) < 1 \forall p \in [0, 1] \iff (p - 1)^2 + p > 0 \forall p \in [0, 1].$$

Let's consider the function of the right inequality when $n \rightarrow \infty$:

$$\lim_{n \rightarrow \infty} (2 + \frac{1}{n})^n (p(1 - p))^{\lfloor \frac{n}{2} \rfloor}$$

We know that this limit is zero because:

$$(2 + \frac{1}{n})^n (p(1 - p))^{\lfloor \frac{n}{2} \rfloor} = \left[(2 + \frac{1}{n})^2 p(1 - p) \right]^{\lfloor \frac{n}{2} \rfloor}$$

and, given $p \in [0, 1]$, always exists a value N such that:

$$\exists N > 0 \text{ s.t. } \forall n > N, \left((2 + \frac{1}{n})^2 p(1 - p) \right) < 1 \iff p(1 - p) < \frac{1}{4 + \frac{1}{n^2} + \frac{4}{n}},$$

which for $n \rightarrow \infty$ holds if and only if $p \neq \frac{1}{2}$.

We can now apply the *squeeze theorem* and show that the derivative is equal to zero if $p \in [0, 1], p \neq \frac{1}{2}$. It is important to notice that $\frac{\partial \mathcal{T}(p)}{\partial p}$ is not continuous in $p = \frac{1}{2}$.

Case 2: n is even.

$$\lim_{n \rightarrow \infty} \frac{\partial \mathcal{T}(p)}{\partial p} = \lim_{n \rightarrow \infty} \frac{1}{p} (p(1-p))^{\frac{n}{2}} \frac{n}{2} \binom{n}{\frac{n}{2}}$$

which is equivalent to:

$$\frac{1}{p} \lim_{n \rightarrow \infty} (p(1-p))^{\frac{n}{2}} \frac{n}{2} \binom{n}{\frac{n}{2}}$$

We saw before that:

$$\lim_{n \rightarrow \infty} (p(1-p))^{\frac{n}{2}} \frac{n}{2} \binom{n}{\frac{n}{2}} = 0$$

Thus, the result holds also for the case in which p is even. □

6.3 Case: $p > 0.5$ or $p < 0.5$

In the previous section, we proved that $\lim_{n \rightarrow \infty} \frac{\partial \mathcal{T}(p)}{\partial p} = 0$ if $p \neq 0.5$. This implies that we can rewrite $\mathcal{T}(p)$ for $n \rightarrow \infty$ in the following form:

$$\lim_{n \rightarrow \infty} \mathcal{T}(p) = \begin{cases} v_1 & p < 0.5 \\ v_2 & p = 0.5 \\ v_3 & p > 0.5, \end{cases} \tag{6}$$

with v_1, v_2 and v_3 real numbers in $[0, 1]$ such that $v_1 \leq v_2 \leq v_3$ (since $\mathcal{T}(p)$ is monotonic). We already proved that $v_2 = 0.5$.

It is easy to see that $v_1 = 0$, because $\mathcal{T}(0) = 0, \forall n$ since all the terms of the sum are equal to zero. Finally, we have that $v_3 = 1$, because $\mathcal{T}(1) = 1, \forall n$.

In fact, $\mathcal{T}(1)$ corresponds to the probability of getting the correct prediction considering a profile of n classifiers where each one elects the correct class with 100% of accuracy. Since we are considering Plurality, which satisfies the axiomatic property of unanimity, the aggregated profile will also elect the correct class with 100% of accuracy. Thus, the value of $\mathcal{T}(1)$ is 1 for each $n > 0$ and consequently for $n \rightarrow \infty$. Thus, we showed that:

$$\lim_{n \rightarrow \infty} \mathcal{T}(p) = \begin{cases} 0 & p < 0.5 \\ 0.5 & p = 0.5 \\ 1 & p > 0.5, \end{cases} \tag{7}$$

This concludes the proof of Theorem 2.

7 Theoretical analysis: relaxing same-accuracy and independence assumptions

In this section we will relax the assumptions made in Sect. 5 in two ways: first, we remove the assumption that each classifier in the profile has the same accuracy p , allowing the classifiers to have a different accuracy (while still considering them independent); later we instead relax the independence assumption, allowing dependencies between classifiers by taking into account the presence of areas of the domain that are correctly classified by at least half of the classifiers simultaneously.

7.1 Independent classifiers with different accuracy values

Considering the same accuracy p for all classifiers is not realistic, even if we set $p = \frac{1}{n} \sum_{i \in A} p_i$, that is, the average profile accuracy. In what follows, we will relax this assumption by extending our study to the general case in which each classifier in the profile can have different accuracy, while still considering them independent. More precisely, we assume that each classifier i has accuracy p_i of choosing the correct class c^* .

In this case the probability of choosing the correct class for our ensemble method is:

$$\frac{1}{K} \sum_{(S_1, \dots, S_m) \in \Omega_{c^*}} \left[\prod_{i \in S^*} (1 - p_i) \cdot \prod_{i \in S^c} p_i \right]$$

where K is the normalization function, S is the set of all classifiers $S = \{1, 2, \dots, n\}$; S_i is the set of classifiers that elect candidate c_i ; S^* is the set of classifiers that elect c^* ; $\overline{S^*}$ is the complement of S^* in S ($\overline{S^*} = S \setminus S^*$); and Ω_{c^*} is the set of all possible partitions of S in which c^* is chosen:

$$\Omega_{c^*} = \{(S_1, \dots, S_{m-1}) \mid \text{partitions of } \overline{S^*} \text{ s.t. } |S_i| < |S^*| \forall i : c_i \neq c^*\}.$$

Notice that this scenario has been analyzed, although from a different point of view, in the literature (see for example [44, 51]). However, the focus of these works is fundamentally different from ours, since their goal is to find the optimal decision rule that maximizes the probability that a profile elects the correct class.

Another relevant work is the one from [38] in which the authors study the case where a profile of n voters have to make a decision over k options. Each voter i has independent probabilities $p_i^1, p_i^2, \dots, p_i^k$ of voting for options 1, 2, \dots , k respectively. The probability, $p_i^{c^*}$ (i.e., the probability of voting for the correct outcome c^*) exceeds each probabilities p_i^c of voting for any of the incorrect outcomes, $c \neq c^*$. The main difference with our approach is that in [38] the authors assume to know the full probability distribution over the outcomes for each voter, moreover they assume the voters have the same probability distribution. In this regard, we just assume to know the accuracy p_i (different for each voter) for each classifier/voter (where $p_i = p_i^{c^*}$). Thus, we provide a more general formula that covers more scenarios.

7.2 Dependent classifiers

Until now, we assumed that the classifiers are independent: the set of the correctly classified examples of a specific classifier is selected by using an independent uniform distribution over all the examples.

We now relax this assumption, by considering dependencies between classifiers by taking into account the presence of areas of the domain that are correctly classified by at least half of the classifiers simultaneously. The idea is to estimate the amount of *overlapping* of the classifications of the individual classifiers. We denote by ρ the ratio of the examples that are in the *easy-to-classify* part of the domain (in which more than half of the classifiers is able to predict the correct label c^*). Thus, ρ equal to 1 when the whole domain is *easy-to-classify*. Considering n classifiers, we can define an upper-bound for ρ :

$$\rho \leq \mathbb{P}[\exists \mathcal{I} \subseteq S, |\mathcal{I}| \geq \frac{n}{2} \text{ s.t. } \forall i \in \mathcal{I} \arg \max(x_i) = c^*].$$

In fact, ρ is bounded by the probability of the correct classification of an example by at least half of the classifiers (which are correctly classified by the ensemble). It is interesting to note that $\rho \leq p$. Removing the *easy-to-classify* examples from the training dataset, we obtain the following accuracy for the other examples:

$$\tilde{p} = \frac{p - \rho}{1 - \rho} < p. \quad (8)$$

We are now ready to generalize Theorem 1.

Theorem 3 *The probability of choosing the correct class c^* in a profile of n classifiers with accuracy $p \in [0, 1]$, m classes and with an overlapping value ρ , using Plurality to compute the winner, is larger than:*

$$(1 - \rho)\mathcal{T}(\tilde{p}) + \rho. \quad (9)$$

The statement follows from Theorem 1 and splitting the correctly classified examples by the ratio defined by ρ . This result tells us that, in order to obtain an improvement of the individual classifiers' accuracy p , we need to maximize the Formula 9. This corresponds to avoid maximizing the overlap ρ (the ratio of the examples that are in the *easy-to-classify* in which more than half of the classifiers is able to predict the correct label) since this would lead to a counter-intuitive effect: if we maximize the overlap of a set of classifiers with accuracy p , in the optimal case the accuracy of the ensemble would be p as well (we recall that ρ is bounded by p). Our goal is instead to obtain a collective accuracy greater than p . Thus, the idea is that we want to focus also on the examples that are more difficult to classify.

The ideal case, to improve the final performance of the ensemble, is to generate a family of classifiers with a balanced trade-off between ρ and the portion of accuracy generated by classifying the difficult examples (i.e., the ones not in the *easy-to-classify* set). A reasonable way to pursue this goal corresponds to choosing the base classifiers randomly.

Example 3 Consider $n = 10$ classifiers with $m = 2$ classes and assume the accuracy of each classifier in the profile is $p = 0.7$. Following the previous observations, we know that $\rho \leq 0.7$. In the case of the maximum overlap among classifiers, i.e., $\rho = 0.7$, the accuracy of VORACE is $0.3\mathcal{T}(\tilde{p}) + 0.7$. Recalling Eq. 8, we have that $\tilde{p} = 0$ and, consequently, $\mathcal{T}(\tilde{p}) = \mathcal{T}(0) = 0$. Thus, the accuracy of VORACE remains exactly 0.7. In general (see Fig. 1), with small values for the input accuracy p , the function $\mathcal{T}(p)$ obtains a decrease of the original accuracy. On the other hand, in the case of a smaller overlap, for example the edge case of $\rho = 0$, we have that $\tilde{p} = p$, and Formula 9 becomes equal to the original Formula 1. Then, VORACE is able to exploit the increase of performance given by $n = 10$

classifiers with a high \tilde{p} of 0.7. In fact, Formula 9 becomes simply $\mathcal{T}(0.7)$ that is close to $0.85 > 0.7$, improving the accuracy of the final model.

8 Conclusions and future work

We have proposed the use of voting rules in the context of ensemble classifiers: a voting rule aggregates the predictions of several randomly generated classifiers, with the goal to obtain a classification that is closer to the correct one. Via a theoretical and experimental analysis, we have shown that this approach generates ensemble classifiers that perform similarly to, or even better than, existing ensemble methods. This is especially true when VORACE employs Plurality or Copeland as voting rules. In particular, Plurality has also the added advantage to require very little information from the individual classifiers and being tractable. Compared to building ad-hoc classifiers that optimize the hyper-parameters configuration for a specific dataset, our approach does not require any knowledge of the domain and thus it is more broadly usable also by non-experts.

We plan to extend our work to deal with other types of data, such as structured data, text, or images. This will also allow for a direct comparison of our approach with the work by [6]. Moreover, we are working on extending the theoretical analysis beyond the Plurality case.

We also plan to consider the extension of our approach to multi-class classification. In this regard, a prominent application of voting theory to this scenario might come from the use of committee selection voting rules [20] in an ensemble classifier. We also plan to study properties of voting rules that may be relevant and desired in the classification domain (see for instance [23, 24]), with the aim to identify and select voting rules that possess such properties, or to define new voting rules with these properties, or also to prove impossibility results about the presence of one or more such properties.

Appendix: Discussion and comparison with [41]

In this section, we compare our theoretical formula to estimate the accuracy of VORACE in Eq. 1 (for the plurality case) with respect to the one provided in [41] (page 93 Sect. 3.2, formula for P_{id} Eq. 8), providing details of the problem of their formulation. From our analysis, we discovered that applying their estimation of the – so called – Identification Rate (P_{id}) produces incorrect results, even in simple cases. We can prove it by using the following counterexample: a binary classification problem where the goal is “to combine” a single classifier with accuracy p , i.e., number of classes $m = 2$, and number of classifiers $n = 1$. It is straightforward that the final accuracy of a combination of a single classifier with accuracy p has to remain unchanged ($P_{id} = p$).

Before proceeding with the calculations, we have to introduce some quantities, following the same ones defined in their original paper:

- N_t is a random variable that gives the total number of votes received by the correct class:

$$P(N_t = j) = \binom{n}{j} p^j (1-p)^{n-j}.$$

- N_s is a random variable that gives the total number of votes received by the wrong class s^{th} :

$$P(N_s = j) = \binom{n}{j} e^j (1 - e)^{n-j},$$

where $e = \frac{1-p}{m-1}$ is the misclassification rate.

- N_s^{max} is a random variable that gives the maximum number of votes among all the wrong classes:

$$\begin{aligned} P(N_s^{max} = k) &= \\ &= \sum_{h=1}^{m-1} \binom{m-1}{h} P(N_s = k)^h P(N_s < j)^{m-1-h}, \end{aligned}$$

where the quantity $P(N_s < j)$ is:

$$P(N_s < j) = \sum_{t=0}^{j-1} P(N_s = t).$$

The authors assume that N_t and N_s^{max} are independent random variables. This means that the probability that the correct class obtains k votes is independent to the probability that the maximum votes within the wrong classes correspond to j . This false assumption leads to a wrong final formula. In fact, applying Eq. 8 in [41] to our simple binary scenario with a single classifier, we have that the new estimated accuracy is:

$$\begin{aligned} P_{id} &= \sum_{j=1}^N P(N_t = j) \sum_{k=0}^{j-1} P(N_s^{max} = k) \\ &= P(N_t = 1) P(N_s^{max} = 0) = p^2, \end{aligned} \tag{10}$$

whereas the correct result should be p .

On the other hand, our proposed formula (Theorem 1) tackles this scenario correctly, as proved in the following, where we specify Equation 1 to this context:

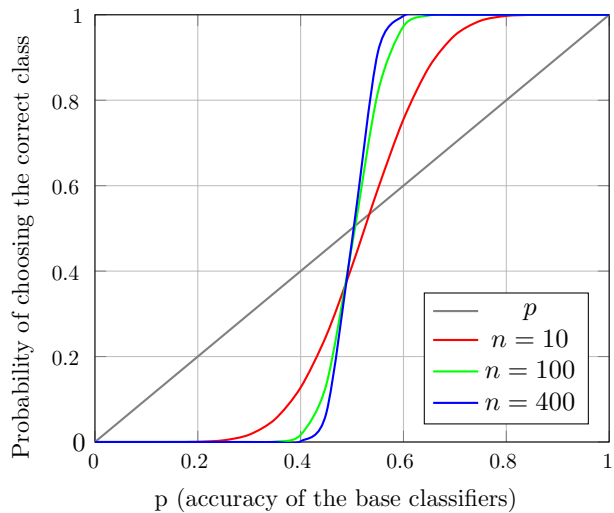
$$\begin{aligned} P_{id} &= \frac{1}{K} (1 - p)^n \sum_{i=\lceil \frac{n}{m} \rceil}^n \varphi_i(n - i)! \binom{n}{i} \left(\frac{p}{1 - p} \right)^i \\ &= \frac{1}{K} \varphi_1(0)! p = p, \end{aligned}$$

where $\varphi_1(0)! = 1$ and $K = 1$.

Notice that, as expected, Formula 1 is equal to 1 when $p = 1$, meaning that, when all classifiers are correct, our ensemble method correctly outputs the same class as all individual classifiers.

As other proof of the difference between the two formulas, we created a similar plot as the one in Fig. 1, applying Eq.(8) in [41]—instead of our formula—obtaining Fig. 3. The two plots are similar, with a less steepness in the curves generated by using our formula. In this sense, we suppose that the formula proposed by [41] is a good approximation of the correct value of P_{id} for large values of n (as we proved that for $n = 1$ and $m = 2$ is not correct).

Fig. 3 Probability of choosing the correct class (P_{id}) varying the size of the profile n in $\{10, 50, 100\}$ and keeping m constant to 2, where each classifier has the same probability p of classifying a given instance correctly, by using Eq. 8 in [41]



Funding Open access funding provided by European University Institute - Fiesole within the CRUI-CARE Agreement. This article has received funding from the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation programme (grant agreement No 833647).

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

References

1. Arrow, K. J., Sen, A. K., & Suzumura, K. (2002). *Handbook of social choice and welfare*. North-Holland.
2. Ateeq, T., Majeed, M. N., Anwar, S. M., Maqsood, M., Rehman, Z., Lee, J. W., et al. (2018). Ensemble-classifiers-assisted detection of cerebral microbleeds in brain MRI. *Computers and Electrical Engineering*, 69, 768–781. <https://doi.org/10.1016/j.compeleceng.2018.02.021>.
3. Azadbakht, M., Fraser, C. S., & Khoshelham, K. (2018). Synergy of sampling techniques and ensemble classifiers for classification of urban environments using full-waveform lidar data. *International Journal of Applied Earth Observation and Geoinformation*, 73, 277–291. <https://doi.org/10.1016/j.jag.2018.06.009>.
4. Barandela, R., Valdovinos, R. M., & Sánchez, J. S. (2003). New applications of ensembles of classifiers. *Pattern Analysis and Applications*, 6(3), 245–256. <https://doi.org/10.1007/s10044-003-0192-z>.
5. Bauer, E., & Kohavi, R. (1999). An empirical comparison of voting classification algorithms: Bagging, boosting, and variants. *Machine Learning*, 36(1–2), 105–139. <https://doi.org/10.1023/A:1007515423169>.
6. Bergstra, J., & Bengio, Y. (2012). Random search for hyper-parameter optimization. *Journal of Machine Learning Research*, 13, 281–305.

7. Breiman, L. (1996a). Bagging predictors. *Machine Learning*, 24(2), 123–140. <https://doi.org/10.1007/BF00058655>.
8. Breiman, L. (1996b). Stacked regressions. *Machine Learning*, 24(1), 49–64.
9. Chen, T., & Guestrin, C. (2016). Xgboost: A scalable tree boosting system. In *Proceedings of the 22nd acm sigkdd international conference on knowledge discovery and data mining, ACM* (pp. 785–794).
10. Condorcet, N. (2014). *Essai sur l'application de l'analyse à la probabilité des décisions rendues à la pluralité des voix (Cambridge Library Collection - Mathematics)*. Cambridge: Cambridge University Press. <https://doi.org/10.1017/CBO9781139923972>
11. Conitzer, V., & Sandholm, T. (2005). Common voting rules as maximum likelihood estimators. In *Proceedings of the Twenty-First Conference on Uncertainty in Artificial Intelligence* (pp. 145–152). Arlington, Virginia, USA: AUAI Press. <http://dl.acm.org/citation.cfm?id=3020336.3020354>
12. Conitzer, V., Davenport, A., & Kalagnanam, J. (2006). Improved bounds for computing kemeny rankings. *AAAI*, 6, 620–626.
13. Conitzer, V., Rognlie, M., & Xia, L. (2009). Preference functions that score rankings and maximum likelihood estimation. In *IJCAI 2009, Proceedings of the 21st international joint conference on artificial intelligence, Pasadena, California, USA, July 11–17, 2009* (pp. 109–115).
14. Cornelio, C., Donini, M., Loreggia, A., Pini, M. S., & Rossi, F. (2020). Voting with random classifiers (VORACE). In *Proceedings of the 19th international conference on autonomous agents and multi-agent systems (AAMAS)* (pp. 1822–1824).
15. De Condorcet, N., et al. (2014). *Essai sur l'application de l'analyse à la probabilité des décisions rendues à la pluralité des voix*. Cambridge: Cambridge University Press.
16. Dietterich, T. G. (2000). An experimental comparison of three methods for constructing ensembles of decision trees: Bagging, boosting, and randomization. *Machine Learning*, 40(2), 139–157.
17. Dietterich, T. G., & Bakiri, G. (1995). Solving multiclass learning problems via error-correcting output codes. *The Journal of Artificial Intelligence Research*, 2, 263–286. <https://doi.org/10.1613/jair.105>.
18. Donini, M., Loreggia, A., Pini, M. S., & Rossi, F. (2018). Voting with random neural networks: A democratic ensemble classifier. In *Proceedings of the RiCeRca Workshop - co-located with the 17th International Conference of the Italian Association for Artificial Intelligence*.
19. van Erp, M., & Schomaker, L. (2000). Variants of the borda count method for combining ranked classifier hypotheses. In *7th workshop on frontiers in handwriting recognition* (pp. 443–452).
20. Faliszewski, P., Skowron, P., Slinko, A., & Talmon, N. Multiwinner voting: A new challenge for social choice theory. *Trends in Computational Social Choice*, 74(2017), 27–47.
21. Freund, Y., & Schapire, R. E. (1997). A decision-theoretic generalization of on-line learning and an application to boosting. *The Journal of Computer and System*, 55(1), 119–139. <https://doi.org/10.1006/jcss.1997.1504>.
22. Gandhi, I., & Pandey, M. (2015). Hybrid ensemble of classifiers using voting. In *2015 international conference on Green Computing and Internet of Things (ICGCIoT)* (pp. 399–404). IEEE.
23. Grandi, U., Loreggia, A., Rossi, F., & Saraswat, V. (2014). From sentiment analysis to preference aggregation. In *International Symposium on Artificial Intelligence and Mathematics, ISAIM*.
24. Grandi, U., Loreggia, A., Rossi, F., & Saraswat, V. (2016). A borda count for collective sentiment analysis. *Annals of Mathematics and Artificial Intelligence*, 77(3), 281–302.
25. Gul, A., Perperoglou, A., Khan, Z., Mahmoud, O., Miftahuddin, M., Adler, W., & Lausen, B. (2018a). Ensemble of a subset of knn classifiers. *Advances Data Analysis and Classification*, 12, 827–840.
26. Gul, A., Perperoglou, A., Khan, Z., Mahmoud, O., Miftahuddin, M., Adler, W., & Lausen, B. (2018b). Ensemble of a subset of knn classifiers. *Advances Data Analysis and Classification*, 12(4), 827–840. <https://doi.org/10.1007/s11634-015-0227-5>.
27. Ho, T. K. (1995). Random decision forests. *Document analysis and recognition, IEEE*, 1, 278–282.
28. Kemeny, J. G. (1959). Mathematics without numbers. *Daedalus*, 88(4), 577–591.
29. Khoshgoftaar, T. M., Hulse, J. V., & Napolitano, A. (2011). Comparing boosting and bagging techniques with noisy and imbalanced data. *IEEE Transactions Systems, Man, and Cybernetics, Part A*, 41(3), 552–568. <https://doi.org/10.1109/TSMCA.2010.2084081>.
30. Kittler, J., Hatef, M., & Duin, R.P.W. (1996). Combining classifiers. In *Proceedings of the Sixth International Conference on Pattern Recognition* (pp. 897–901). Silver Spring, MD: IEEE computer society press.
31. Kotsiantis, S. B., & Pintelas, P. E. (2005). Local voting of weak classifiers. *KES Journal*, 9(3), 239–248.
32. Kotsiantis, S. B., Zaharakis, I. D., & Pintelas, P. E. (2006). Machine learning: A review of classification and combining techniques. *Artificial Intelligence Review*, 26(3), 159–190. <https://doi.org/10.1007/s10462-007-9052-3>.

33. Kuncheva, L., Whitaker, C., Shipp, C., & Duin, R. (2003). Limits on the majority vote accuracy in classifier fusion. *Pattern Analysis and Applications*, 6(1), 22–31. <https://doi.org/10.1007/s10044-002-0173-7>.
34. Lam, L., & Suen, S. (1997). Application of majority voting to pattern recognition: An analysis of its behavior and performance. *IEEE Transactions on Systems, Man, and Cybernetics*, 27, 553–567.
35. Leon, F., Floria, S. A., & Badica, C. (2017). Evaluating the effect of voting methods on ensemble-based classification. In *INISTA-17*, (pp. 1–6). <https://doi.org/10.1109/INISTA.2017.8001122>
36. Leung, K. T., & Parker, D. S. (2003). Empirical comparisons of various voting methods in bagging. *Proceedings of the Ninth ACM SIGKDD International Conference on Knowledge Discovery and Data Mining* (pp. 595–600). NY, USA: ACM.
37. Yacoub, S., Lin, X., & Simske, S. (2003). Performance analysis of pattern classifier combination by plurality voting. *Pattern Recognition Letters*, 24, 1959–1969.
38. List, C., & Goodin, R. (2001). Epistemic democracy: Generalizing the condorcet jury theorem. *Journal of Political Philosophy*. <https://doi.org/10.1111/1467-9760.00128>.
39. Loreggia, A., Mattei, N., Rossi, F., & K. Brent Venable. (2018). Preferences and Ethical Principles in Decision Making. In *Proceedings of the 2018 AAAI/ACM Conference on AI, Ethics, and Society (AI/ES '18)*. New York, NY: Association for Computing Machinery. <https://doi.org/10.1145/3278721.3278723>
40. Melville, P., Shah, N., Mihalkova, L., & Mooney, R. J. (2004). Experiments on ensembles with missing and noisy data. In *Multiple Classifier Systems, 5th International Workshop, MCS 2004, Cagliari, Italy, June 9–11, 2004* (pp. 293–302). https://doi.org/10.1007/978-3-540-25966-4_29
41. Mu, X., Watta, P., & Hassoun, M. H. (2009). Analysis of a plurality voting-based combination of classifiers. *Neural Processing Letters*, 29(2), 89–107. <https://doi.org/10.1007/s11063-009-9097-1>.
42. Neto, A. F., & Canuto, A. M. P. (2018). An exploratory study of mono and multi-objective metaheuristics to ensemble of classifiers. *Applied Intelligence*, 48(2), 416–431. <https://doi.org/10.1007/s10489-017-0982-4>.
43. Newman, C. B. D., & Merz, C. (1998). *UCI repository of machine learning databases*. <http://www.ics.uci.edu/~mllearn/MLRepository.html>
44. Nitzan, S., & Paroush, J. (1982). Optimal decision rules in uncertain dichotomous choice situations. *International Economic Review*, 23(2), 289–297.
45. Perikos, I., & Hatzilygeroudis, I. (2016). Recognizing emotions in text using ensemble of classifiers. *Engineering Applications of Artificial Intelligence*, 51, 191–201.
46. Rokach, L. (2010). Ensemble-based classifiers. *Artificial Intelligence Review*, 33(1–2), 1–39.
47. Rossi, F., & Loreggia, A. (2019). Preferences and ethical priorities: Thinking fast and slow in AI. In *Proceedings of the 18th Autonomous Agents and Multi-agent Systems Conference* (pp. 3–4).
48. Rossi, F., Venable, K. B., & Walsh, T. (2011). A short introduction to preferences: Between artificial intelligence and social choice. *Synthesis lectures on artificial intelligence and machine learning, Morgan & claypool publishers*. <https://doi.org/10.2200/S00372ED1V01Y201107AIM014>.
49. Saleh, E., Blaszczyński, J., Moreno, A., Valls, A., Romero-Aroca, P., de la Riva-Fernandez, S., & Slowinski, R. (2018). Learning ensemble classifiers for diabetic retinopathy assessment. *Artificial Intelligence in Medicine*, 85, 50–63. <https://doi.org/10.1016/j.artmed.2017.09.006>.
50. Moulin, H. (2016). In F. Brandt, V. Conitzer, U. Endriss, J. Lang, & A. Procaccia (Eds.), *Handbook of Computational Social Choice*. Cambridge: Cambridge University Press. <https://doi.org/10.1017/CBO9781107446984>
51. Shapley, L., & Grofman, B. (1984). Optimizing group judgmental accuracy in the presence of interdependencies. *Public Choice*.
52. Strubell, E., Ganesh, A., & McCallum, A. (2019). Energy and policy considerations for deep learning in NLP. In *Proceedings of the 57th conference of the association for computational linguistics, ACL 2019, Florence, Italy, July 28- August 2, 2019*, Vol. 1 (pp. 3645–3650). Long Papers.
53. Sun, X., Lin, X., Shen, S., & Hu, Z. (2017). High-resolution remote sensing data classification over urban areas using random forest ensemble and fully connected conditional random field. *ISPRS International Journal of Geo-Information*, 6(8), 245. <https://doi.org/10.3390/ijgi6080245>.
54. Webb, G. I. (2000). Multiboosting: A technique for combining boosting and wagging. *Machine Learning*, 40(2), 159–196. <https://doi.org/10.1023/A:1007659514849>.
55. Young, H. (1988). Condorcet's theory of voting. *American Political Science Review*, 82(4), 1231–1244. <https://doi.org/10.2307/1961757>