



Confidence sets with expected sizes for Multiclass Classification

Christophe Denis, Mohamed Hebiri

► To cite this version:

Christophe Denis, Mohamed Hebiri. Confidence sets with expected sizes for Multiclass Classification. 2016. hal-01357850v1

HAL Id: hal-01357850

<https://hal.science/hal-01357850v1>

Preprint submitted on 30 Aug 2016 (v1), last revised 28 Nov 2017 (v2)

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Confidence sets with expected sizes for Multiclass Classification

Christophe Denis* and Mohamed Hebiri†

LAMA, UMR-CNRS 8050,
Université Paris Est – Marne-la-Vallée

Abstract

Challenging multiclass classification problems such as image annotation may involve a large number of classes. In this context, confusion between classes may occur, and single label classification may be misleading. We provide in the present paper a general device that, given a classification procedure and an unlabeled dataset, outputs a set of class labels, instead of a single one. Interestingly, this procedure does not require that the unlabeled dataset explores the whole classes. Even more, the method is calibrated to control the expected size of the output set while minimizing the classification risk. We show the statistical optimality of the procedure and establish rates of convergence under the Tsybakov margin condition. It turns out that these rates are linear on the number of labels. We illustrate the numerical performance of the procedure on simulated and on real data. In particular, we show that with moderate expected size, w.r.t. the number of labels, the procedure provides significant improvement of the classification risk.

Keywords : Multiclass classification, confidence sets, plug-in confidence sets, cumulative distribution functions.

1 Introduction

Multiclass classification consists in assigning a single label $Y \in \{1, \dots, K\}$ to a given input example $X \in \mathcal{X}$ among a collection of labels. In particular, in the last decay, high-dimensional statistics have witnessed the rise of problems where the number of classes is large, such as in face recognition or locations tagging. While large number of classes provide precise characterization of an instance, they also might induce high confusion between classes. That is to say, many examples would belong to the margin between several classes. In this context, classification procedures often have weak accuracy and then relevant classification may involve more than a single label as output. However, we keep in mind that an informative output would rather contain only a few labels. As an interesting application, we can consider image annotations problems where the number of classes is large (*e.g.* 1000) and the goal is then to pick a small number (*e.g.* 10) of candidate labels. Given an instance X , a natural approach would be to consider the class labels associated to the the top level conditional probabilities $\mathbb{P}(Y = \cdot | X = x)$. In the present paper, we show that such strategy is suboptimal; we provide an alternative approach.

The scope of the present paper is then to build confidence sets for multiclass classification problems for which we are able to control the expected size. A confidence set Γ is a function

*Christophe.Denis@u-pem.fr

†Mohamed.Hebiri@u-pem.fr

that maps \mathcal{X} onto $2^{\mathcal{Y}}$ and the expected size of Γ is defined as $\mathbb{E}[|\Gamma(X)|]$, where $|\cdot|$ stands for the cardinality. More precisely, for an example X , given an expected set size β and consistent estimators of the conditional probabilities, the algorithm outputs a set $\hat{\Gamma}(X)$ such that $\mathbb{E}[|\hat{\Gamma}(X)|] \approx \beta$. Furthermore, the procedure aims at minimizing the classification error, that is:

$$\mathbb{P}(Y \notin \hat{\Gamma}(X)) \approx \min_{\Gamma: \mathbb{E}[|\Gamma(X)|] = \beta} \mathbb{P}(Y \notin \Gamma(X)) = \mathcal{R}_{\beta}^*.$$

We exhibit an oracle confidence set, says Γ^* , which exactly achieves under mild assumption the risk \mathcal{R}_{β}^* for a fixed β . The definition of the oracle confidence set Γ^* relies on the cumulative distribution functions of the conditional probabilities. Therefore, our procedure is based on the plug-in rule. We prove that our algorithm performs asymptotically as well as the oracle procedure. Moreover, we derive rates of convergence under the Tsybakov noise assumption on the data generating distribution. In particular, we show that this rate is linear on K , the number of classes. Up to our knowledge, the above described method as well as its statistical study have not been studied before in the multiclass setting. According to numerical considerations, our procedure is implementable in a semi-supervised way [Vap98]: labeled data are only used for the estimation of the conditional probabilities. This step is performed with any machine learning algorithm such as deep learning or random forest procedures. In addition, we do not need that the unlabelled examples explore the whole classes.

Related works: The closest learning task to the present work is *classification with reject option* which is a particular setting in binary classification. Several papers fall within the scope of this area [Cho70, HW06, WY11, Lei14, DH15] and differ from each other by the goal they consider. Our procedure is inspired by the paper [DH15] that also considers a semi-supervised approach to build confidence sets invoking cumulative distribution functions of the conditional probabilities. However, most of the aspects developed in [DH15] are not directly adaptable to the multiclass case. In particular, we develop here new probabilistic results on sums of cumulative distribution functions of random variables, that are of own interest.

Assigning a set of labels instead of a single one for an input example is not new [VGS05, WLW04, dCDB09, LRW13, CCB16]. One of the most popular methods is based on *Conformal Prediction* approach [VGS99, Vov02, VGS05]. In the multiclass classification framework, the goal of this algorithm is to build the smallest set of labels such that its classification error is below a pre-specified level. Since our procedure aims at minimizing the classification error while keeping under control the size of the set, *Conformal Prediction* can be seen as a dual of our method. We refer to the very interesting statistical study of *Conformal Predictors* in the binary case in the paper [Lei14].

Notation: First, we state general notation. Let $\mathcal{Y} = \{1, \dots, K\}$, with $K \geq 2$ being an integer. Let (X, Y) be the generic data-structure taking its values in $\mathcal{X} \times \mathcal{Y}$ with distribution \mathbb{P} . The goal in classification is to predict the label Y given an observation of X . This is performed based on a classifier (or classification rule) s which is a function mapping \mathcal{X} onto \mathcal{Y} . Let \mathcal{S} be the set of all classifiers. The misclassification risk R associated with $s \in \mathcal{S}$ is defined as

$$R(s) = \mathbb{P}(s(X) \neq Y).$$

Moreover, the minimizer of R over \mathcal{S} is the Bayes classifier, denoted by s^* , and is characterized by

$$s^*(\cdot) = \operatorname{argmax}_{k \in \mathcal{Y}} p_k(\cdot),$$

where $p_k(x) = \mathbb{P}(Y = k | X = x)$ for $x \in \mathcal{X}$ and $k \in \mathcal{Y}$.

Let us now consider more specific notation related to the multiclass confidence set setting. Let

a confidence set be any measurable function that maps \mathcal{X} onto $2^{\mathcal{Y}}$. Let Γ be a confidence set. This confidence set is characterized by two attributes. The first one is the risk associated to the confidence set

$$\mathcal{R}(\Gamma) = \mathbb{P}(Y \notin \Gamma(X)), \quad (1)$$

and is related to its accuracy. The second attribute is linked to the information given by the confidence set. It is defined as

$$\mathcal{I}(\Gamma) = \mathbb{E}(|\Gamma(X)|), \quad (2)$$

where $|\cdot|$ stands for the cardinality. Moreover, for some $\beta \in [1, K]$, we say that, for two confidence sets Γ and Γ' such that $\mathcal{I}(\Gamma) = \mathcal{I}(\Gamma') = \beta$, the confidence set Γ is “better” than Γ' if $\mathcal{R}(\Gamma) \leq \mathcal{R}(\Gamma')$.

Organization of the paper: The rest of the paper is organized as follows. Next section is devoted to the definition and the main properties of the oracle confidence set for multiclass classification. The plug-in counterpart of this oracle, as well as its asymptotic behavior are provided in Section 3. Rates of convergence are also stated in that section. We present a numerical study of our methods on simulated and real data in Section 4. We finally draw some conclusions and present perspectives of our work in Section 5. Proofs of our results are postponed to the Appendix.

2 Confidence set for multiclass classification

In the present section, we define a class of confidence sets adapted to multiclass classification referred as *Oracle β -sets*. For some $\beta \in (0, K)$, these sets are shown to be optimal according to the risk (1) with an information (2) equal to β . Moreover, basic but fundamental properties of Oracle β -sets can be found in Proposition 1, while Proposition 3 provides another interpretation of these sets.

2.1 Notation and definition

First of all, we introduce in this section a class of confidence sets that specifies oracle confidence sets. Let $\beta \in (0, K)$ be a desired information level. The so-called *Oracle β -sets* are optimal according to the risk (1) among all the confidence sets Γ such that $\mathcal{I}(\Gamma) = \beta$. Throughout the paper we make the following assumption

(A1) For all $k \in \{1, \dots, K\}$, the cumulative distribution function F_{p_k} of $p_k(X)$ is continuous.

The definition of the *Oracle β -set* relies on the continuous decreasing function G defined for any $t \in [0, 1]$ by

$$G(t) = \sum_{k=1}^K \bar{F}_{p_k}(t),$$

where for any $k \in \{1, \dots, K\}$, we denote by \bar{F}_{p_k} the tail distribution function of $p_k(X)$, that is, $\bar{F}_{p_k} = 1 - F_{p_k}$ with F_{p_k} being the cumulative distribution function (c.d.f.) of $p_k(X)$. We define the generalized inverse G^{-1} of G (see [vdV98]):

$$G^{-1}(\beta) = \inf\{t \in [0, 1] : G(t) \leq \beta\}, \quad \forall \beta \in (0, K).$$

The functions G and G^{-1} are central in the construction of the Oracle β -sets. We then provide some of their useful properties in the following proposition.

Proposition 1. *The following properties on G hold*

- i) *For every $t \in (0, 1)$ and $\beta \in (0, K)$, $G^{-1}(\beta) \leq t \Leftrightarrow \beta \geq G(t)$.*
- ii) *For every $\beta \in (0, K)$, $G(G^{-1}(\beta)) = \beta$.*
- iii) *Let ε be a random variable, independent of X , distributed from a uniform distribution on $\{1, \dots, K\}$ and let U be distributed from a uniform distribution on $[0, K]$. Define*

$$Z = \sum_{k=1}^K p_k(X) \mathbf{1}_{\{\varepsilon=k\}}.$$

If the function G is continuous, then $G(Z) \stackrel{L}{=} U$ and $G^{-1}(U) \stackrel{L}{=} Z$.

Now, we are able to define the Oracle β -set:

Definition 1. *Let $\beta \in (0, K)$, the Oracle β -set is given by*

$$\begin{aligned} \Gamma_{\beta}^*(X) &= \{k \in \{1, \dots, K\} : G(p_k(X)) \leq \beta\} \\ &= \{k \in \{1, \dots, K\} : p_k(X) \geq G^{-1}(\beta)\}. \end{aligned}$$

This definition of the Oracle β -set is intuitive and closely related to the binary classification with reject option setting [Cho70, HW06, DH15]: the label k is assigned to the Oracle β -set if the probability $p_k(X)$ is large enough. It is worth noting that the function G plays the same role as the c.d.f. of the score function in [DH15]. As provided in Proposition 1, their introduction allows to control exactly the information (2). Indeed, it follows from the definition of the Oracle β -set that for each $\beta \in (0, K)$

$$|\Gamma_{\beta}^*(X)| = \sum_{k=1}^K \mathbf{1}_{\{p_k(X) \geq G^{-1}(\beta)\}},$$

and then $\mathcal{I}(\Gamma_{\beta}^*) = G(G^{-1}(\beta))$. Therefore, Proposition 1 ensures that

$$\mathcal{I}(\Gamma_{\beta}^*) = \beta.$$

This last display points out that the Oracle β -sets are indeed β -level according to the information parameter. In the next section, we focus on the study of the risk of these oracle confidence sets.

Remark 1. *Naturally, the definition of Oracle β -sets can be extended to any $\beta \in [0, K]$. However, the limit cases $\beta = 0$ and $\beta = K$ are of limited interest and are completely trivial. We then exclude these two limit cases from the present study.*

2.2 Properties of the oracle confidence sets

Let us first state the optimality of the Oracle β -set:

Proposition 2. *Let Assumption (A1) be satisfied. For any $\beta \in (0, K)$, we have both:*

1. *The Oracle β -set Γ_{β}^* satisfies the following property:*

$$\mathcal{R}(\Gamma_{\beta}^*) = \inf_{\Gamma} \mathcal{R}(\Gamma),$$

where the infimum is taken over all β -level confidence sets.

2. For any β -level confidence set Γ , the following holds

$$0 \leq \mathcal{R}(\Gamma) - \mathcal{R}(\Gamma_\beta^*) = \mathbb{E} \left[\sum_{k \in (\Gamma_\beta^*(X) \Delta \Gamma(X))} |p_k(X) - G^{-1}(\beta)| \right], \quad (3)$$

where the symbol Δ stands for the symmetric difference of two sets, that is, for two subsets A and B of $\{1, \dots, K\}$, we write $A \Delta B = (A \setminus B) \cup (B \setminus A)$.

Several remarks can be made from Proposition 2. First, for $\beta \in (0, K)$, the Oracle β confidence set is optimal, according to its risk, over the class of β -sets. Moreover, the excess risk of any β -set relies on the behavior of the score functions p_k around the threshold $G^{-1}(\beta)$. Finally, we can note that if $K = 2$ and $\beta = 1$ which implies $G^{-1}(\beta) = 1/2$, Equation (3) reduced to the misclassification excess risk in binary classification.

Remark 2. An important motivation behind the introduction of confidence sets and in particular of Oracle β -set is that they might outperform the Bayes rule which can be seen as the Oracle β -set associated to $\beta = 1$. This gap in performance is even larger when the number of classes K is large and there is a big confusion between classes. Such improvement will be illustrated in the numerical study.

Next, we provide another characterization of the Oracle β -set.

Proposition 3. For $t \in [0, 1]$, and Γ a confidence set, we define

$$L_t(\Gamma) = \mathbb{P}(Y \notin \Gamma(X)) + t\mathcal{I}(\Gamma).$$

For $\beta \in [0, K]$, the following equality holds:

$$L_{G^{-1}(\beta)}(\Gamma_\beta^*) = \min_{\Gamma} L_{G^{-1}(\beta)}(\Gamma).$$

The proof of this proposition relies on the same arguments as in Proposition 2. It is then omitted. Proposition 3 states that the Oracle β -set is defined as the minimizer, over all confidence sets Γ , of the risk function L_t when the tuning parameter t is set equal to $G^{-1}(\beta)$. Note moreover that the risk function L_t is a trade-off, controlled by the parameter t , of the risk of a confidence set on the one hand, and of the information provided by this confidence set on the other hand. Hence, the risk function L_t can be viewed as a generalization to the multiclass case of the risk function provided in [Cho70, HW06] for binary classification with reject option setting. The above characterization of the Oracle β -set will be of particular interest when dealing with empirical risk minimization which will be the scope future works. In the present present, we consider instead a plug-in approach which is detailed in the next section.

3 Plug-in β -sets

This section is devoted to the study of the empirical counterpart of the Oracle β -set defined by plug-in rule. The construction of these confidence sets, referred as *Plug-in β -sets*, is given in Section 3.1. Several theoretical properties, such as asymptotic consistency and rates of convergence are set in Section 3.2.

3.1 Data driven procedure

In this section we provide a simple data driven procedure which mimics the Oracle β set. This procedure is a generalization of the plug-in method in [DH15] provided for the binary case. Our procedure relies on preliminary estimators of the functions p_k , $k = 1, \dots, K$. To set these estimators, we introduce a first data set \mathcal{D}_n , which consists of independent copies of (X, Y) . Thereafter, for each $k \in \{1, \dots, K\}$, we denote \hat{p}_k the estimator of p_k computed based on \mathcal{D}_n . Moreover, conditional on \mathcal{D}_n , we introduce the c.d.f. of $\hat{p}_k(X)$

$$F_{\hat{p}_k}(t) = \mathbb{P}_X(\hat{p}_k(X) \leq t | \mathcal{D}_n),$$

for $t \in [0, 1]$. As for the c.d.f. of p_k , we make the following assumption:

(A2) *The cumulative distribution functions $F_{\hat{p}_k}$ with $k = 1, \dots, K$ are continuous.*

Remark 3. *As will be seen in the numerical study, even in cases where Assumption (A2) does not hold, one can provide a regularized version of the procedure that implies that the above assumption always hold.*

At this point, we are able to define a plug-in approximation of the Oracle β -set for $\beta \in (0, K)$:

$$\tilde{\Gamma}_\beta(X) = \left\{ k \in \{1, \dots, K\} : \tilde{G}(\hat{p}_k(X)) \leq \beta \right\}, \quad (4)$$

where for $t \in [0, 1]$

$$\tilde{G}(t) = \sum_{k=1}^K \bar{F}_{\hat{p}_k}(t),$$

with $\bar{F}_{\hat{p}_k} = 1 - F_{\hat{p}_k}$. Since the function \tilde{G} depends only on the unknown distribution of X , it is enough to consider a second *unlabeled* dataset $\mathcal{D}_N = \{X_i, i = 1, \dots, N\}$, independent of \mathcal{D}_n in order to compute the empirical versions of the $\bar{F}_{\hat{p}_k}$'s. Finally, we consider the plug-in β -sets

Definition 2. *Let $\beta \in (0, K)$ and \hat{p}_k be any estimators of p_k , $k = 1, \dots, K$. The plug-in β -set is defined as follows:*

$$\hat{\Gamma}_\beta(X) = \left\{ k \in \{1, \dots, K\} : \hat{G}(\hat{p}_k(X)) \leq \beta \right\},$$

where

$$\hat{G}(\cdot) = \frac{1}{N} \sum_{i=1}^N \sum_{k=1}^K \mathbf{1}_{\{\hat{p}_k(X_i) \geq \cdot\}}.$$

The most important remark about the construction of this data driven procedure is that it is made in a semi-supervised way. Indeed, the estimation of the tail distribution functions of \hat{p}_k requires only a set of *unlabeled* observations. This is particularly attractive in applications where the number of label observations is small (because labelling examples is time consuming or may be costly) and where one has in hand several unlabeled features that can be used to make prediction more accurate. As an important consequence is that the estimation of the tail distribution functions of \hat{p}_k does not depend on the number of observations in each class label. That is to say, this unlabeled dataset can be unbalanced with respect to the classes, that can often occur in multiclass classification settings where the number of classes K is quite large. Next section deals with the theoretical performance of the plug-in β -set.

3.2 Theoretical performance

This section is devoted to the study of asymptotic properties of the plug-in β -sets. First, we state some additional notation. The symbols \mathbf{P} and \mathbf{E} stand for generic probability and expectation, respectively. Moreover, given a plug-in β -set $\hat{\Gamma}_\beta$, we consider its risk $\mathbf{R}(\hat{\Gamma}_\beta) = \mathbf{P}(Y \notin \hat{\Gamma}_\beta(X))$ and its information $\mathbf{I}(\hat{\Gamma}_\beta) = \mathbf{E}[|\hat{\Gamma}_\beta(X)|]$. We also write $a_n = O(b_n)$ for two non-negative sequences (a_n) and (b_n) if there exists some non-negative constant $c > 0$ such that: $a_n \leq c b_n$ for n sufficiently large. Let us now state our main results.

Proposition 4. *For each $\beta \in (0, K)$, the following equalities hold*

1. $\mathbf{E}[|\tilde{\Gamma}_\beta(X)|] = \beta$,
2. $\mathbf{I}(\hat{\Gamma}_\beta) = \beta + O\left(\frac{K}{\sqrt{N}}\right)$ and
3. $\mathbf{E}[|\hat{\Gamma}_\beta(X) \Delta \tilde{\Gamma}_\beta(X)|] = O\left(\frac{K}{\sqrt{N}}\right)$.

The above result illustrates the behavior of the plug-in β -set with respect to the semi-oracle approximation $\tilde{\Gamma}_\beta$ defined by (4). Proposition 4 says that $\tilde{\Gamma}_\beta$ reaches the exact level of information β , while the plug-in β -set $\hat{\Gamma}_\beta$ is of level β up to a term of order K/\sqrt{N} . Last point of the proposition even says that $\hat{\Gamma}_\beta$ and $\tilde{\Gamma}_\beta$ are asymptotically the same sets.

Next, exploiting the result in the above proposition, theorem makes the comparison between the plug-in β -set and the Oracle β -set given in Definition 1.

Theorem 1. *1. If $\hat{p}_k(X) \rightarrow p_k(X)$ in probability when $n \rightarrow +\infty$ for every $k \in \{1, \dots, K\}$, then for any $\beta \in (0, K)$,*

$$\begin{aligned} \mathbf{R}(\hat{\Gamma}_\beta) &\rightarrow \mathcal{R}(\Gamma_\beta^*), \\ \mathbf{E}[|\hat{\Gamma}_\beta(X) \Delta \Gamma_\beta^*(X)|] &\rightarrow 0 \end{aligned}$$

when both n and N go to infinity.

2. *For any $\beta \in (0, K)$, assume that there exist $C_1 < \infty$ and $\gamma_\beta > 0$ such that for every k*

$$\mathbb{P}(|p_k(X) - G^{-1}(\beta)| \leq t) \leq C_1 t^{\gamma_\beta}, \quad \forall t > 0. \quad (5)$$

Assume also that there exist a sequence of positive numbers $a_n \rightarrow +\infty$ and some positive constants C_2, C_3 such that for every k

$$\mathbf{P}(|\hat{p}_k(x) - p_k(x)| \geq t) \leq C_2 \exp(-C_3 a_n t^2), \quad \forall t > 0, \forall x \in \mathcal{X}. \quad (6)$$

Then we have

$$\mathbf{R}(\hat{\Gamma}_\beta) - \mathcal{R}(\Gamma_\beta^*) = K \left(O(a_n^{-\gamma_\beta/2}) + O(N^{-1/2}) \right) \quad (7)$$

$$\mathbf{E}[|\hat{\Gamma}_\beta(X) \Delta \Gamma_\beta^*(X)|] = K \left(O(a_n^{-\gamma_\beta/2}) + O(N^{-1/2}) \right). \quad (8)$$

Recall that throughout the paper, Assumptions (A1) and (A2) are supposed to be satisfied. Under these two mild assumptions, the above result states that the Oracle β -sets mimics asymptotically the Oracle β -set in terms of its risk and its information. The second part of this theorem deals with rates of convergence. These rates are linear in K , the number of classes that seems to be the characteristic of multiclass classification as compared to binary classification. As a consequence, the bounds we get generalize the bounds for confidence sets in the binary classification setting obtained in [DH15]. In particular, (7) and (8) requires classical assumptions in classification as the Tsybakov margin condition (5) with a slight adaptation to the multiclass setting [Tsy04, AT07, DH15, HW06]. This assumption, together with the condition (6) are needed only to state the first part of the rate related to the estimation of the conditional probabilities p_k . We refer the reader to the paper [DH15] where a longer discussion on these conditions is provided. Also, both of these conditions are not used to provide the second part of the rates which is of order K/\sqrt{N} and relies on the estimation of the function G . Hence, this part of the estimation is established under the mild Assumptions (A1) and (A2).

4 Numerical results

In this section, we evaluate the plug-in β -set numerically. First, in Section 4.1, a simulation study is led to support our theory. Plug-in β -set are then applied in Section 4.2 to real dataset.

4.1 Simulation study

We perform a simulation study in order to evaluate the performance of our plug-in procedure. We compare our method with the strategy which is based on the β top levels conditional probabilities. In the sequel this method is referred as **max** procedure. We also investigate the influence of the parameter K . Hence, we consider two simulation schemes relying on the parameter $K \in \{10, 100\}$. For each K , we generate (X, Y) according to a mixture model. More precisely,

- i) the label Y is distributed from a uniform distribution on $\{1, \dots, K\}$;
- ii) conditional on $Y = k$, the feature X is generated according to a multivariate gaussian distribution with mean parameter $\mu_k \in \mathbb{R}^{10}$ and identity covariance matrix.

For each $k = 1, \dots, K$, the vectors μ_k are i.i.d realizations of uniform distribution on $[0, 4]$. Note that for this distribution, we have

$$p_k(X) = \frac{f_k(X)}{\sum_{j=1}^K f_j(X)},$$

where for $k = 1, \dots, K$, $f_k(X)$ is the density function of a multivariate gaussian distribution with mean parameter μ_k and identity covariance matrix. For each K , the misclassification risk of Bayes rule is evaluated based on sufficiently large dataset. It is valued at 0.22 and at 0.60 for $K = 10$ and for $K = 100$ respectively. Hence, we can observe the confusion induced by the large number classes. In the sequel, we aim at providing the estimation of the risk of the Oracle β -set. To do so, for $\beta \in \{2, 5, 10, 20\}$ and each K , we repeat B times the following steps.

- i) simulate two datasets \mathcal{D}_N and \mathcal{D}_M with $N = 10000$ and $M = 1000$;
- ii) based on \mathcal{D}_N , we compute the empirical counterpart of G and provide an approximation of the Oracle β -set Γ_β^* given in Definition 1 (we recall that this step requires a dataset which contains only unlabeled features);

$K = 10$		
β	Oracle β -set	max
2	0.05 (0.01)	0.09 (0.01)
5	0.00 (0.00)	0.01 (0.00)

$K = 100$		
β	Oracle β -set	max
2	0.39 (0.01)	0.42 (0.01)
5	0.20 (0.01)	0.22 (0.01)
10	0.09 (0.01)	0.11 (0.01)
20	0.03 (0.01)	0.04 (0.01)

Table 1: For each of the $B = 100$ repetitions and each model, we derive the estimated risks \mathcal{R}_M of the Oracle β -set and of the **max** procedure w.r.t. β . We compute the means and standard deviations (between parentheses) over the $B = 100$ repetitions. Top: the data are generated according to $K = 10$ – Bottom: the data are generated according to $K = 100$.

β	$K = 10$	$K = 100$
2	2.00 (0.03)	2.00 (0.03)
5	5.00 (0.08)	5.00 (0.06)
10	.	10.00 (0.13)
20	.	20.02 (0.31)

Table 2: For each of the $B = 100$ repetitions and each model, we derive the estimated informations \mathcal{I}_M of the Oracle β -set w.r.t. β . We compute the means and standard deviations (between parentheses) over the $B = 100$ repetitions. Left: the data are generated according to $K = 10$ – Right: the data are generated according to $K = 100$.

iii) finally, over \mathcal{D}_M , we compute the empirical counterparts \mathcal{R}_M (of $\mathcal{R}(\Gamma_\beta^*)$) and \mathcal{I}_M (of $\mathcal{I}(\Gamma_\beta^*)$).

From this estimates, we compute mean and standard deviation of \mathcal{R}_M and \mathcal{I}_M . The results are provided in Tables 1 and 2. Now for each K and each β , we evaluate the performances of three plug-in β sets. The estimation of the p_k 's is based on random forests, softmax regression and deep learning procedures. We notify that for random forests and softmax regression algorithms, the random variables $\hat{p}_k(X)$ appears to be not continuous. Hence Assumption (A2) is not always satisfied. Therefore, we add to $\hat{p}_k(X)$ an independent small perturbation $|\mathcal{N}(0, 1e^{-10})|$. The evaluation of these plug-in β sets rely on the following steps

- i) simulate three datasets \mathcal{D}_n , \mathcal{D}_N and \mathcal{D}_M ;
- ii) based on \mathcal{D}_n , we compute the estimators \hat{p}_k of p_k according to the considered procedure;
- iii) based on \mathcal{D}_N and \hat{p}_k we compute the function \hat{G} and the plug-in β -set $\hat{\Gamma}_\beta$ as in Definition 2 (we recall that this step requires a dataset which contains only unlabeled features);
- iv) finally, we compute over \mathcal{D}_M the empirical counterpart of \mathbf{R} and of \mathbf{I} of the considered plug-in β -set.

From these experiments, we compute means and standard deviations. For $K = 10$, we fix $n = 1000$ and $N \in \{100, 10000\}$; for $K = 100$ we fix $n = 10000$ and $N \in \{100, 10000\}$. Finally, we fix $M = 1000$. The results are provided in Table 3 and 4. As a reference point for the sequel of our experiments, the misclassification risks of the random forests, softmax regression and deep

$K = 10$								
β -set				max				
β		rforest	softmax reg	deep learn		rforest	softmax reg	deep learn
2		0.09 (0.01)	0.06 (0.01)	0.09 (0.01)		0.13 (0.01)	0.10 (0.01)	0.13 (0.02)
5		0.01 (0.00)	0.00 (0.00)	0.01 (0.00)		0.02 (0.00)	0.01 (0.00)	0.02 (0.00)

$K = 100$								
β -set				max				
β		rforest	softmax reg	deep learn		rforest	softmax reg	deep learn
2		0.48 (0.02)	0.93 (0.01)	0.46 (0.02)		0.51 (0.01)	0.96 (0.01)	0.48 (0.02)
5		0.30 (0.02)	0.85 (0.02)	0.25 (0.02)		0.31 (0.01)	0.90 (0.01)	0.27 (0.01)
10		0.17 (0.01)	0.75 (0.02)	0.12 (0.01)		0.18 (0.01)	0.80 (0.01)	0.14 (0.01)
20		0.07 (0.01)	0.59 (0.02)	0.04 (0.01)		0.09 (0.01)	0.61 (0.02)	0.06 (0.01)

Table 3: For each of the $B = 100$ repetitions and for each model, we derive the estimated risks \mathbf{R} of three different plug-in β -sets w.r.t. β . We compute the means and standard deviations (between parentheses) over the $B = 100$ repetitions. For each β and for each N , the plug-in β -sets, as well as the **max** procedures are based on, from left to right, **rforest**, **softmax reg** and **deep learn**, which are respectively the random forest, the softmax regression and the deep learning methods. Top: the data are generated according to $K = 10$ – Bottom: the data are generated according to $K = 100$.

$K = 10$								
$N = 100$				$N = 10000$				
β		rforest	softmax reg	deep learn		rforest	softmax reg	deep learn
2		2.01 (0.09)	2.01 (0.10)	2.02 (0.11)		2.00 (0.02)	2.00 (0.03)	2.00 (0.03)
5		5.02 (0.18)	4.99 (0.20)	5.00 (0.21)		5.00 (0.06)	5.00 (0.08)	5.00 (0.07)

$K = 100$								
$N = 100$				$N = 10000$				
β		rforest	softmax reg	deep learn		rforest	softmax reg	deep learn
2		2.02 (0.10)	2.09 (0.43)	2.01 (0.09)		2.00 (0.03)	2.02 (0.15)	2.00 (0.02)
5		4.97 (0.15)	5.27 (0.70)	5.01 (0.24)		5.00 (0.04)	5.01 (0.27)	5.00 (0.07)
10		9.98 (0.24)	10.02 (1.00)	10.02 (0.42)		10.01 (0.09)	10.05 (0.32)	10.00 (0.16)
20		20.08 (0.48)	19.74 (0.98)	20.11 (0.85)		20.00 (0.16)	20.01 (0.36)	20.01 (0.28)

Table 4: For each of the $B = 100$ repetitions and for each model, we derive the estimated informations \mathbf{I} of three different plug-in β -sets w.r.t. β and the sample size N . We compute the means and standard deviations (between parentheses) over the $B = 100$ repetitions. For each β and each N , the plug-in β -sets are based on, from left to right, **rforest**, **softmax reg** and **deep learn**, which are respectively the random forest, the softmax regression and the deep learning procedures. Top: the data are generated according to $K = 10$ – Bottom: the data are generated according to $K = 100$.

learning methods are valued respectively at 0.28, 0.24, 0.29 for $K = 10$, and at 0.65, 0.98 0.63 for $K = 100$.

First, from Tables 2 and 4, we observe that the approximation of the information is good. The approximation accuracy increases with N the number of non-labeled data. Let us now consider the experiment results related to the classification error. Table 1 teaches us that the performances of the Oracle β -set get better with the value of the parameter β . Although, for moderate β as compared to K we obtain satisfactory improvement. For instance when $K = 10$ and $\beta = 2$ the classification error rate decreases from 0.22 to 0.05; likewise, when $K = 100$ and $\beta = 5$ the the classification error rate decreases from 0.60 to 0.20. Second, Table 3 provide same conclusions regarding the plug-in rules. However, we note that for $K = 100$ the convergence of the plug-in rules is a bit slower except for softmax regression procedure which has poor performances. Finally, Tables 1 and 3 show that the **max** procedure is slightly outperformed by our method. This was expected, but the **max** procedure does not require the unlabeled dataset.

4.2 Application on real dataset

In this section, we provide an application of the plug-in β -set on two real datasets: the *Forest type mapping* dataset and the *one-hundred plant species leaves* dataset coming from the UCI database. In the sequel we refer to these two datasets as **Forest** and **Plant** respectively. The **Forest** dataset consists of $K = 4$ classes and 523 labeled observations (we gather the train and test sets) with 27 features. Here the classes are unbalanced. In the **Plant** dataset, there are 100 classes and 1600 labeled observations. This dataset is balanced so that each class consists of 16 observations. The original dataset contains 3 covariates (each covariate consists of 64 features). In order to make the problem more challenging, we drop 2 covariates.

For these two datasets, we evaluate risks (and informations) of the plug-in β -set based on the same estimation procedures (*i.e.*, the random forest, the softmax regression and the deep learning procedures). To do so, we use the *cross validation* principle. In particular, we run $B = 100$ times the procedure where we split the data each time in three: a sample of size n for the estimation of the conditional distributions p_k 's; a sample of size N to estimate the function G ; and a sample of size M to evaluate the risk and the information. For both datasets, we make sure that in the sample of size n , there is the same number of observations in each class. Moreover, let notify that the misclassification risks of the random forest, the softmax regression and the deep learning procedures in the **Forest** dataset are respectively 0.14, 0.20 and 0.26, whereas in the **Plant** dataset, they are evaluated to 0.38, 0.92 and 0.82 respectively. As planned, the performance of the classical methods are quite bad in this last dataset. In particular, both softmax regression and deep learning procedures perform poorly.

As in the simulation section, we compare the plug-in β -set to the **max** procedure. The size of the samples are set to be equal to $n = 200$, $N = 100$ and $M = 223$ for the **Forest** dataset, and $n = 1000$, $N = 200$ and $M = 400$ for the **Plant** one. The results, reported in Table 5, confirm the observations made in Section 4.1. In particular, we see that the classification risk gets drastically better with moderate β . For instance, for the **Plant**, and with $\beta = 10$ the classification error rate of the plug-in β -set based on deep learning procedure decreases from 0.82 to 0.27.

5 Conclusion

In multiclass classification setting, the present paper propose a new procedure that assigns a set of labels instead of a single label to each instance. This set has a controlled expected size (or information) and its construction relies on cumulative distribution functions. In this sense the proposed methods is a generalization of the works in [DH15] in the binary classification with reject

Forest ($K = 4$)									
β -set					max				
β			rforest	softmax reg	deep learn		rforest	softmax reg	deep learn
2	R		0.02 (0.01)	0.06 (0.03)	0.06 (0.03)		0.03 (0.01)	0.06 (0.02)	0.07 (0.03)
	I		2.01 (0.09)	2.02 (0.08)	2.01 (0.09)		.	.	.
Plant ($K = 100$)									
β -set					max				
β			rforest	softmax reg	deep learn		rforest	softmax reg	deep learn
2	R		0.21 (0.02)	0.77 (0.01)	0.71 (0.02)		0.24 (0.01)	0.84 (0.02)	0.72 (0.02)
	I		2.00 (0.08)	2.03 (0.13)	2.00 (0.07)		.	.	.
10	R		0.03 (0.01)	0.43 (0.02)	0.27 (0.03)		0.05 (0.01)	0.44 (0.02)	0.25 (0.02)
	I		9.95 (0.24)	9.97 (0.29)	10.02 (0.32)		.	.	.

Table 5: For each of the $B = 100$ repetitions and for each dataset, we derive the estimated risks **R** and information **I** of three different plug-in β -sets w.r.t. β . We compute the means and standard deviations (between parentheses) over the $B = 100$ repetitions. For each β , the plug-in β -sets, as well as the **max** procedures are based on, from left to right, **rforest**, **softmax reg** and **deep learn**, which are respectively the random forest, the softmax regression and the deep learning methods. Top: the dataset is the **Forest** – the dataset is the **Plant**.

option. Theoretical guarantees, especially rates of convergence are also provided and involve the continuity of these cumulative distribution functions. These bounds differ from those established in the binary case by a factor K , which is the number of classes (see [DH15] for the comparison). However, techniques of proof relies on different arguments and new probabilistic results (see Lemma 1 below). The plug-in β -set procedure described in Section 3 is defined as a two steps algorithm whose second step consists in the estimation of the function G (which is a sum of tail distribution functions). Interestingly, this step does not require a set of labeled data and neither to explore the whole classes, that is suitable for semi-supervised learning. Future works will focus on optimality in the minimax sense. In particular, we will investigate whether the rates of convergence are optimal in terms of their dependence on K . Based on the characterization of the Oracle β -sets provided in Proposition 3, we also intent to study procedures based on empirical risk minimization [YW10].

6 Appendix

This section gathers the proofs of our results.

6.1 Technical Lemmas

We first lay out key lemmata, which are crucial to establish the main theory. Note that Lemma 2 is inspired by Lemma 3.1 in [AT07]. We consider $K \geq 2$ be an integer, and Z_1, \dots, Z_K , K random variables. Moreover we define function H by:

$$H(t) = \frac{1}{K} \sum_{k=1}^K F_k(t), \quad \forall t \in [0, 1],$$

where for all $k = 1, \dots, K$, F_k is the cumulative distribution function of Z_k . Finally, let us define the generalized inverse H^{-1} of H :

$$H^{-1}(p) = \inf\{t : H(t) \geq p\}, \quad \forall p \in (0, 1).$$

Lemma 1. *Let ε distributed from a uniform distribution on $\{1, \dots, K\}$ and independent of Z_k , $k = 1, \dots, K$. Let U distributed from a uniform distribution on $[0, 1]$. We consider*

$$Z = \sum_{k=1}^K Z_k \mathbf{1}_{\{\varepsilon=k\}}.$$

If H is continuous then

$$H(Z) \stackrel{\mathcal{L}}{=} U \quad \text{and} \quad H^{-1}(U) \stackrel{\mathcal{L}}{=} Z$$

Proof. First we note that for every $t \in [0, 1]$, $\mathbb{P}(H(Z) \leq t) = \mathbb{P}(Z \leq H^{-1}(t))$. Moreover, we have

$$\begin{aligned} \mathbb{P}(Z \leq t) &= \sum_{k=1}^K \mathbb{P}(Z \leq H^{-1}(t), \varepsilon = k) \\ &= \frac{1}{K} \sum_{k=1}^K \mathbb{P}(Z_k \geq H^{-1}(t)) \quad \text{with } \varepsilon \text{ independent of } X \\ &= H(H^{-1}(t)) \\ &= t \quad \text{with } H \text{ continuous.} \end{aligned}$$

To conclude the proof, we observe that

$$\begin{aligned} \mathbb{P}(H^{-1}(U) \leq t) &= \mathbb{P}(U \geq H(t)) \\ &= \frac{1}{K} \sum_{k=1}^K F_k(t) \\ &= \sum_{k=1}^K \mathbb{P}(Z_k \leq t, \varepsilon = k) \\ &= \mathbb{P}(Z \leq t). \end{aligned}$$

□

Lemma 2. *Let X be a real random variable, $(X_n)_{n \geq 1}$ a sequence of real random variables and $t_0 \in \mathbb{R}$. Assume that there exist $C_1 < \infty$ and $\gamma_0 > 0$ such that*

$$\mathbb{P}_X(|X - t_0| \leq \delta) \leq C_1 \delta^{\gamma_0}, \quad \forall \delta > 0,$$

and a sequence of positive numbers $a_n \rightarrow +\infty$, C_2, C_3 some positive constants such that

$$\mathbb{P}_{X_n}(|X_n - X| \geq \delta | X) \leq C_2 \exp(-C_3 a_n \delta^2), \quad \forall \delta > 0, \quad \forall n \in \mathbb{N}.$$

Then, there exists $C > 0$ depending only on C_1, C_2 and C_3 , such that

$$\begin{aligned} |\mathbf{E}[\mathbf{1}_{\{X_n \geq t_0\}} - \mathbf{1}_{\{X \geq t_0\}}]| &\leq \mathbf{E}[|\mathbf{1}_{\{X_n \geq t_0\}} - \mathbf{1}_{\{X \geq t_0\}}|] \\ &\leq \mathbf{P}(|X_n - X| \geq |X - t_0|) \\ &\leq C a_n^{-\gamma_0/2}. \end{aligned}$$

Lemma 3 (Hoeffding Inequality). *Let $b > 0$ be a real number, and N be a positive integer. Let X_1, \dots, X_N be N random variables strictly bounded by the intervals $[0, b]$, then*

$$\mathbb{P} \left(\left| \frac{1}{N} \sum_{i=1}^N (X_i - \mathbb{E}[X_i]) \right| \geq t \right) \leq 2 \exp \left(-\frac{2Nt^2}{b^2} \right), \quad \forall t > 0.$$

6.2 Proof of Proposition 2

Let $\beta > 0$ and Γ a confidence set such that $\mathcal{I}(\Gamma) = \beta$. First, we note that the following decomposition holds

$$\begin{aligned} \mathcal{R}(\Gamma) - \mathcal{R}(\Gamma_\beta^*) &= \sum_{j=1}^K \sum_{k=1}^K \mathbb{E} \left[\sum_{l=1}^K (\mathbf{1}_{\{Y \notin \Gamma(X)\}} - \mathbf{1}_{\{Y \notin \Gamma_\beta^*(X)\}}) \mathbf{1}_{\{Y=l\}} \mathbf{1}_{\{|\Gamma(X)|=k\}} \mathbf{1}_{\{|\Gamma_\beta^*(X)|=j\}} \right] \\ &= \sum_{j=1}^K \sum_{k=1}^K \mathbb{E} \left[\sum_{l=1}^K (\mathbf{1}_{\{l \notin \Gamma(X)\}} - \mathbf{1}_{\{l \notin \Gamma_\beta^*(X)\}}) p_l(X) \mathbf{1}_{\{|\Gamma(X)|=k\}} \mathbf{1}_{\{|\Gamma_\beta^*(X)|=j\}} \right] \\ &= \sum_{j=1}^K \sum_{k=1}^K \mathbb{E} \left[\sum_{l=1}^K (\mathbf{1}_{\{l \in \Gamma_\beta^*(X) \setminus \Gamma(X)\}} - \mathbf{1}_{\{l \in \Gamma(X) \setminus \Gamma_\beta^*(X)\}}) p_l(X) \mathbf{1}_{\{|\Gamma(X)|=k, |\Gamma_\beta^*(X)|=j\}} \right]. \end{aligned}$$

From the last decomposition, with

$$\begin{aligned} \mathbb{E}[|\Gamma(X)|] &= \mathbb{E}[|\Gamma_\beta^*(X)|] \\ &= \mathbb{E} \left[\sum_{j=1}^K \sum_{k=1}^K k \mathbf{1}_{\{|\Gamma(X)|=k\}} \mathbf{1}_{\{|\Gamma_\beta^*(X)|=j\}} \right] \\ &= \mathbb{E} \left[\sum_{j=1}^K \sum_{k=1}^K j \mathbf{1}_{\{|\Gamma(X)|=k\}} \mathbf{1}_{\{|\Gamma_\beta^*(X)|=j\}} \right], \end{aligned}$$

we can express the excess risk as the sum of two terms:

$$\begin{aligned} \mathcal{R}(\Gamma) - \mathcal{R}(\Gamma_\beta^*) &= \sum_{j=1}^K \sum_{k=1}^K \mathbb{E} \left[\left(\sum_{l=1}^K \mathbf{1}_{\{l \in \Gamma_\beta^*(X) \setminus \Gamma(X)\}} p_l(X) - j G^{-1}(\beta) \right) \mathbf{1}_{\{|\Gamma(X)|=k\}} \mathbf{1}_{\{|\Gamma_\beta^*(X)|=j\}} \right] \\ &\quad + \sum_{j=1}^K \sum_{k=1}^K \mathbb{E} \left[\left(k G^{-1}(\beta) - \sum_{l=1}^K \mathbf{1}_{\{l \in \Gamma(X) \setminus \Gamma_\beta^*(X)\}} p_l(X) \right) \mathbf{1}_{\{|\Gamma(X)|=k\}} \mathbf{1}_{\{|\Gamma_\beta^*(X)|=j\}} \right]. \quad (9) \end{aligned}$$

Now, for $j, k \in \{1, \dots, K\}$ on the event $\{|\Gamma(X)| = k, |\Gamma_\beta^*(X)| = j\}$, we have

$$k = \sum_{l=1}^K \mathbf{1}_{\{l \in \Gamma(X) \setminus \Gamma_\beta^*(X)\}} + \sum_{l=1}^K \mathbf{1}_{\{l \in \Gamma(X) \cap \Gamma_\beta^*(X)\}},$$

and

$$j = \sum_{l=1}^K \mathbf{1}_{\{l \in \Gamma_\beta^*(X) \setminus \Gamma(X)\}} + \sum_{l=1}^K \mathbf{1}_{\{l \in \Gamma(X) \cap \Gamma_\beta^*(X)\}}.$$

Therefore, since

$$l \in \Gamma_\beta^* \Leftrightarrow p_l(X) \geq G^{-1}(\beta),$$

Equality (9) yields the result.

6.3 Proof of Asymptotic results

We first set a Lemma that will be used in the proof of our result. This lemma relies on some arguments given in [AT07] and Lemma 1.

6.3.1 Tool lemma

Lemma 4. *There exists an absolute constant $C > 0$ such that*

$$\sum_{k=1}^K \mathbf{P} \left(|\hat{G}(\hat{p}_k(X)) - \tilde{G}(\hat{p}_k(X))| \geq |\tilde{G}(\hat{p}_k(X)) - \beta| \right) \leq \frac{CK}{\sqrt{N}}.$$

Proof. We define, for $\delta > 0$ and $k \in \{1, \dots, K\}$

$$\begin{aligned} \mathbf{A}_0^k &= \left\{ \left| \tilde{G}(\hat{p}_k(X)) - \beta \right| \leq \delta \right\} \\ \mathbf{A}_j^k &= \left\{ 2^{j-1}\delta < |\tilde{G}(\hat{p}_k(X)) - \beta| \leq 2^j\delta \right\}, \quad j \geq 1. \end{aligned}$$

Since, for every k , the events $(\mathbf{A}_j^k)_{j \geq 0}$ are mutually exclusive, we deduce

$$\begin{aligned} \sum_{k=1}^K \mathbf{P} \left(|\hat{G}(\hat{p}_k(X)) - \tilde{G}(\hat{p}_k(X))| \geq |\tilde{G}(\hat{p}_k(X)) - \beta| \right) &= \\ \sum_{k=1}^K \sum_{j \geq 0} \mathbf{P} \left(|\hat{G}(\hat{p}_k(X)) - \tilde{G}(\hat{p}_k(X))| \geq |\tilde{G}(\hat{p}_k(X)) - \beta|, \mathbf{A}_j^k \right). \end{aligned} \quad (10)$$

Now, we consider ε uniformly distributed on $\{1, \dots, K\}$ independent of \mathcal{D}_n and X . Conditional on \mathcal{D}_n and under Assumption (A_2) , we apply Lemma 1 with $Z_k = \hat{p}_k(X)$, $Z = \sum_{k=1}^K Z_k \mathbf{1}_{\{\varepsilon=k\}}$ and then obtain that $\tilde{G}(Z)$ is uniformly distributed on $[0, K]$. Therefore, for all $j \geq 0$ and $\delta > 0$, we deduce

$$\begin{aligned} \frac{1}{K} \sum_{k=1}^K \mathbb{P}_X \left(|\tilde{G}(\hat{p}_k(X)) - \beta| \leq 2^j\delta | \mathcal{D}_n \right) &= \mathbb{P}_X \left(|\tilde{G}(Z) - \beta| \leq 2^j\delta | \mathcal{D}_n \right) \\ &\leq \frac{2^{j+1}\delta}{K}. \end{aligned}$$

Hence, for all $j \geq 0$, we obtain

$$\sum_{k=1}^K \mathbf{P}(\mathbf{A}_j^k) \leq 2^{j+1}\delta. \quad (11)$$

Next, we observe that for all $j \geq 1$

$$\begin{aligned} \sum_{k=1}^K \mathbf{P} \left(|\hat{G}(\hat{p}_k(X)) - \tilde{G}(\hat{p}_k(X))| \geq |\tilde{G}(\hat{p}_k(X)) - \beta|, \mathbf{A}_j^k \right) &\leq \\ \sum_{k=1}^K \mathbb{E}_{(\mathcal{D}_n, X)} \left[\mathbb{P}_{\mathcal{D}_n} \left(|\hat{G}(\hat{p}_k(X)) - \tilde{G}(\hat{p}_k(X))| \geq 2^{j-1}\delta | \mathcal{D}_n, X \right) \mathbf{1}_{\mathbf{A}_j^k} \right]. \end{aligned} \quad (12)$$

Now, since conditional on (\mathcal{D}_n, X) , $\hat{G}(\hat{p}_k(X))$ is an empirical mean of i.i.d random variables of common mean $\tilde{G}(\hat{p}_k(X)) \in [0, K]$, we deduce from Hoeffding's inequality that

$$\mathbb{P}_{\mathcal{D}_N} \left(|\hat{G}(\hat{p}_k(X)) - \tilde{G}(\hat{p}_k(X))| \geq 2^{j-1} \delta | \mathcal{D}_n, X \right) \leq 2 \exp \left(-\frac{N \delta^2 2^{2j-1}}{K^2} \right).$$

Therefore, from Inequalities (10), (11) and (12), we get

$$\begin{aligned} \sum_{k=1}^K \mathbf{P} \left(|\hat{G}(\hat{p}_k(X)) - \tilde{G}(\hat{p}_k(X))| \geq |\tilde{G}(\hat{p}_k(X)) - \beta| \right) \\ \leq \sum_{k=1}^K \mathbf{P}(A_0^k) + \sum_{j \geq 1} 2 \exp \left(-\frac{N \delta^2 2^{2j-1}}{K^2} \right) \left(\sum_{k=1}^K \mathbf{P}(A_j^k) \right) \\ \leq 2\delta + \delta \sum_{j \geq 1} 2^{j+2} \exp \left(-\frac{N \delta^2 2^{2j-1}}{K^2} \right). \end{aligned}$$

Finally, choosing $\delta = \frac{K}{\sqrt{N}}$ in the above inequality, we finish the proof of the lemma. \square

6.3.2 Proof of Proposition 4

1. for every $\beta \in (0, K)$, conditional on \mathcal{D}_n we define

$$\tilde{G}^{-1}(\beta) = \inf\{t \in [0, 1] : \tilde{G}(t) \leq \beta\}. \quad (13)$$

We note that Assumption (A2) ensures that $t \mapsto \tilde{G}(t)$ is continuous and then

$$\tilde{G}(\tilde{G}^{-1}(\beta)) = \beta.$$

Now, we have

$$\begin{aligned} |\tilde{\Gamma}_\beta(X)| &= \sum_{k=1}^K \mathbf{1}_{\{\tilde{G}(\hat{p}_k(X)) \leq \beta\}} \\ &= \sum_{k=1}^K \mathbf{1}_{\{\hat{p}_k(X) \geq \tilde{G}^{-1}(\beta)\}}. \end{aligned}$$

Hence, the last equation implies that

$$\mathbb{E}_X \left[|\tilde{\Gamma}_\beta(X)| | \mathcal{D}_n \right] = \sum_{k=1}^K \mathbb{P}_X \left(\hat{p}_k(X) \geq \tilde{G}^{-1}(\beta) | \mathcal{D}_n \right) = \tilde{G}(\tilde{G}^{-1}(\beta)) = \beta. \quad (14)$$

Therefore, we obtain $\mathbf{E} \left[|\tilde{\Gamma}_\beta(X)| \right] = \beta$.

2. and 3. From 1), we can write

$$\begin{aligned}
\left| \mathbf{E} \left[\left| \hat{\Gamma}_\beta(X) \right| \right] - \beta \right| &\leq \left| \mathbf{E} \left[\left| \hat{\Gamma}_\beta(X) \right| - \left| \tilde{\Gamma}_\beta(X) \right| \right] \right| \\
&\leq \left| \mathbf{E} \left[\sum_{k=1}^K \left(\mathbf{1}_{\{\hat{G}(\hat{p}_k(X)) \leq \beta\}} - \mathbf{1}_{\{\tilde{G}(\hat{p}_k(X)) \leq \beta\}} \right) \right] \right| \\
&\leq \mathbf{E} \left[\left| \hat{\Gamma}_\beta(X) \Delta \tilde{\Gamma}_\beta(X) \right| \right] \\
&\leq \sum_{k=1}^K \mathbf{E} \left[\left| \mathbf{1}_{\{\hat{G}(\hat{p}_k(X)) \leq \beta\}} - \mathbf{1}_{\{\tilde{G}(\hat{p}_k(X)) \leq \beta\}} \right| \right] \\
&\leq \sum_{k=1}^K \mathbf{P} \left(\left| \hat{G}(\hat{p}_k(X)) - \tilde{G}(\hat{p}_k(X)) \right| \geq \left| \tilde{G}(\hat{p}_k(X)) - \beta \right| \right).
\end{aligned}$$

Hence, applying Lemma 4 in the above inequality, we obtain the desired result.

6.3.3 Proof of Theorem 1

We first study the convergence of $\mathbf{R}(\hat{\Gamma}_\beta) - \mathcal{R}(\Gamma_\beta^*)$. According to the excess risk, we first write the decomposition

$$\mathbf{R}(\hat{\Gamma}_\beta) - \mathcal{R}(\Gamma_\beta^*) = \left(\mathbf{R}(\hat{\Gamma}_\beta) - \mathbf{R}(\tilde{\Gamma}_\beta) \right) + \left(\mathbf{R}(\tilde{\Gamma}_\beta) - \mathcal{R}(\Gamma_\beta^*) \right), \quad (15)$$

and deal with each term in the r.h.s. separately.

First, we note that

$$\mathbf{1}_{\{Y \notin \hat{\Gamma}_\beta(X)\}} - \mathbf{1}_{\{Y \notin \tilde{\Gamma}_\beta(X)\}} = \sum_{k=1}^K \mathbf{1}_{\{Y=k\}} \mathbf{1}_{\{k \notin \hat{\Gamma}_\beta(X)\}} - \sum_{k=1}^K \mathbf{1}_{\{Y=k\}} \mathbf{1}_{\{k \notin \tilde{\Gamma}_\beta(X)\}}.$$

Therefore, we can write

$$\begin{aligned}
\mathbf{E} \left[\mathbf{1}_{Y \notin \hat{\Gamma}_\beta(X)} - \mathbf{1}_{Y \notin \tilde{\Gamma}_\beta(X)} \right] &= \sum_{k=1}^K \mathbf{E} \left[p_k(X) \left(\mathbf{1}_{\{k \notin \hat{\Gamma}_\beta(X)\}} - \mathbf{1}_{\{k \notin \tilde{\Gamma}_\beta(X)\}} \right) \right] \\
&= \sum_{k=1}^K \mathbf{E} \left[p_k(X) \left(\mathbf{1}_{\{\hat{G}(\hat{p}_k(X)) > \beta\}} - \mathbf{1}_{\{\tilde{G}(\hat{p}_k(X)) > \beta\}} \right) \right].
\end{aligned}$$

Since $0 \leq p_k(X) \leq 1$ for all $k \in \{1, \dots, K\}$, the last equality implies

$$\begin{aligned}
\left| \mathbf{R}(\hat{\Gamma}_\beta) - \mathbf{R}(\tilde{\Gamma}_\beta) \right| &= \left| \mathbf{E} \left[\mathbf{1}_{Y \notin \hat{\Gamma}_\beta(X)} - \mathbf{1}_{Y \notin \tilde{\Gamma}_\beta(X)} \right] \right| \\
&\leq \sum_{k=1}^K \mathbf{E} \left[\left| \mathbf{1}_{\{\hat{G}(\hat{p}_k(X)) > \beta\}} - \mathbf{1}_{\{\tilde{G}(\hat{p}_k(X)) > \beta\}} \right| \right] \\
&\leq \sum_{k=1}^K \mathbf{P} \left(\left| \hat{G}(\hat{p}_k(X)) - \tilde{G}(\hat{p}_k(X)) \right| \geq \left| \tilde{G}(\hat{p}_k(X)) - \beta \right| \right).
\end{aligned}$$

Therefore, Lemma 4 implies

$$\left| \mathbf{R}(\hat{\Gamma}_\beta) - \mathbf{R}(\tilde{\Gamma}_\beta) \right| \leq \frac{CK}{\sqrt{N}}. \quad (16)$$

Next, we study the second term in the r.h.s (15).

First, we have shown in the proof of Proposition 4 that $\mathbb{E}_X [\tilde{\Gamma}_\beta(X) | \mathcal{D}_n] = \beta$ (see Equation (14)). Hence, Proposition 2 implies that

$$\mathbb{E}_X \left[\left(\mathbf{1}_{\{Y \notin \tilde{\Gamma}_\beta(X)\}} - \mathbf{1}_{\{Y \notin \Gamma_\beta^*(X)\}} \right) | \mathcal{D}_n \right] = \mathbb{E}_X \left[\sum_{k \in (\Gamma_\beta^*(X) \Delta \tilde{\Gamma}_\beta(X))} |p_k(X) - G^{-1}(\beta)| | \mathcal{D}_n \right]. \quad (17)$$

For the sequel, we refer to Equation (13) for the definition of $\tilde{G}^{-1}(\beta)$. Now, we observe that if $k \in (\tilde{\Gamma}_\beta(X) \setminus \Gamma_\beta^*(X))$, then:

- on the event $\{\tilde{G}^{-1}(\beta) > G^{-1}(\beta)\}$, we have $|\hat{p}_k(X) - p_k(X)| \geq |p_k(X) - G^{-1}(\beta)|$;
- on the event $\{\tilde{G}^{-1}(\beta) \leq G^{-1}(\beta)\}$ either $\hat{p}_k(X) \in (\tilde{G}^{-1}(\beta), G^{-1}(\beta))$
or $|\hat{p}_k(X) - p_k(X)| \geq |p_k(X) - G^{-1}(\beta)|$.

Similar reasoning holds if $k \in (\Gamma_\beta^*(X) \setminus \tilde{\Gamma}_\beta(X))$. Therefore, Equation (17) yields

$$\begin{aligned} & \mathbb{E}_X \left[\left(\mathbf{1}_{\{Y \notin \tilde{\Gamma}_\beta(X)\}} - \mathbf{1}_{\{Y \notin \Gamma_\beta^*(X)\}} \right) | \mathcal{D}_n \right] \\ & \leq \sum_{k=1}^K \mathbb{E}_X \left[|p_k(X) - G^{-1}(\beta)| \mathbf{1}_{\{|\hat{p}_k(X) - p_k(X)| \geq |p_k(X) - G^{-1}(\beta)|\}} | \mathcal{D}_n \right] \\ & \quad + \mathbf{1}_{\{\tilde{G}^{-1}(\beta) \leq G^{-1}(\beta)\}} \sum_{k=1}^K \mathbb{E}_X \left[\mathbf{1}_{\{\hat{p}_k(X) \in (\tilde{G}^{-1}(\beta), G^{-1}(\beta))\}} | \mathcal{D}_n \right] \\ & \quad + \mathbf{1}_{\{G^{-1}(\beta) < \tilde{G}^{-1}(\beta)\}} \sum_{k=1}^K \mathbb{E}_X \left[\mathbf{1}_{\{\hat{p}_k(X) \in (G^{-1}(\beta), \tilde{G}^{-1}(\beta))\}} | \mathcal{D}_n \right]. \quad (18) \end{aligned}$$

Since $\tilde{G}(\tilde{G}^{-1}(\beta)) = G(G^{-1}(\beta))$, we have

$$\begin{aligned} & \mathbf{1}_{\{\tilde{G}^{-1}(\beta) \leq G^{-1}(\beta)\}} \sum_{k=1}^K \mathbb{E}_X \left[\mathbf{1}_{\{\hat{p}_k(X) \in (\tilde{G}^{-1}(\beta), G^{-1}(\beta))\}} | \mathcal{D}_n \right] \\ & + \mathbf{1}_{\{G^{-1}(\beta) \leq \tilde{G}^{-1}(\beta)\}} \sum_{k=1}^K \mathbb{E}_X \left[\mathbf{1}_{\{\hat{p}_k(X) \in (G^{-1}(\beta), \tilde{G}^{-1}(\beta))\}} | \mathcal{D}_n \right] = \left| \tilde{G}(\tilde{G}^{-1}(\beta)) - \tilde{G}(G^{-1}(\beta)) \right| \\ & = \left| G(G^{-1}(\beta)) - \tilde{G}(G^{-1}(\beta)) \right|. \quad (19) \end{aligned}$$

Hence, from Equation (18) and (19), we deduce

$$\begin{aligned} \mathbf{R}(\tilde{\Gamma}_\beta) - \mathcal{R}(\Gamma_\beta^*) & \leq \sum_{k=1}^K \mathbf{E} \left[|p_k(X) - G^{-1}(\beta)| \mathbf{1}_{\{|\hat{p}_k(X) - p_k(X)| \geq |p_k(X) - G^{-1}(\beta)|\}} \right] \\ & \quad + \mathbf{E} \left[\left| G(G^{-1}(\beta)) - \tilde{G}(G^{-1}(\beta)) \right| \right]. \quad (20) \end{aligned}$$

Now, by definition of \tilde{G} , we have

$$\begin{aligned} \mathbf{E} \left[\left| G(G^{-1}(\beta)) - \tilde{G}(G^{-1}(\beta)) \right| \right] \\ \leq \sum_{k=1}^K \mathbf{E} \left[\left| \mathbf{1}_{\{\hat{p}_k(X) > G^{-1}(\beta)\}} - \mathbf{1}_{\{p_k(X) > G^{-1}(\beta)\}} \right| \right] \\ \leq \sum_{k=1}^K \mathbf{P} \left(|\hat{p}_k(X) - p_k(X)| \geq |p_k(X) - G^{-1}(\beta)| \right). \end{aligned}$$

Therefore, with the above Inequality and Equation (20) we have

$$0 \leq \mathbf{R}(\tilde{\Gamma}_\beta) - \mathcal{R}(\Gamma_\beta^*) \leq 2 \sum_{k=1}^K \mathbf{P} \left(|\hat{p}_k(X) - p_k(X)| \geq |p_k(X) - G^{-1}(\beta)| \right). \quad (21)$$

Now, we study $\mathbf{E}(|\hat{\Gamma}_\beta(X) \Delta \Gamma_\beta^*(X)|)$. First, we note that

$$|\hat{\Gamma}_\beta(X) \Delta \Gamma_\beta^*(X)| \leq |\hat{\Gamma}_\beta(X) \Delta \tilde{\Gamma}_\beta(X)| + |\tilde{\Gamma}_\beta(X) \Delta \Gamma_\beta^*(X)|. \quad (22)$$

Proposition 4 ensures that

$$\mathbf{E} \left[|\hat{\Gamma}_\beta(X) \Delta \tilde{\Gamma}_\beta(X)| \right] = O \left(\frac{K}{\sqrt{N}} \right). \quad (23)$$

Moreover, we note that

$$|\tilde{\Gamma}_\beta(X) \Delta \Gamma_\beta^*(X)| = \sum_{k \in (\Gamma_\beta^* \Delta \tilde{\Gamma}_\beta)} 1,$$

and therefore, as for Equation (17), we deduce that

$$\mathbf{E} \left[|\tilde{\Gamma}_\beta(X) \Delta \Gamma_\beta^*(X)| \right] \leq 2 \sum_{k=1}^K \mathbf{P} \left(|\hat{p}_k(X) - p_k(X)| \geq |p_k(X) - G^{-1}(\beta)| \right). \quad (24)$$

Finally, for $0 < \delta_n \rightarrow 0$, we have

$$\begin{aligned} 2 \sum_{k=1}^K \mathbf{P} \left(|\hat{p}_k(X) - p_k(X)| \geq |p_k(X) - G^{-1}(\beta)| \right) \\ \leq 2 \sum_{k=1}^K \mathbf{P} \left(|\hat{p}_k(X) - p_k(X)| \geq \delta_n \right) + 2 \sum_{k=1}^K \mathbf{P} \left(|p_k(X) - G^{-1}(\beta)| \leq \delta_n \right). \end{aligned} \quad (25)$$

So, if $\hat{p}_k(X) \rightarrow p_k(X)$ in probability, we have $\mathbf{P}(|\hat{p}_k(X) - p_k(X)| \geq \delta_n) \rightarrow 0$ for each k . Moreover, Assumption (A1) ensures that $\mathbf{P}(|p_k(X) - G^{-1}(\beta)| \leq \delta_n) \rightarrow 0$ for each k . Therefore, from Inequalities (21), (24) and (25) we obtain

$$\mathbf{R}(\hat{\Gamma}_\beta) - \mathcal{R}(\Gamma_\beta^*) \rightarrow 0 \quad \text{and} \quad \mathbf{E} \left[|\hat{\Gamma}_\beta(X) \Delta \Gamma_\beta^*(X)| \right] \rightarrow 0.$$

For rates of convergence, we apply Lemma 2 in Inequalities (21) and (24).

References

- [AT07] J.Y. Audibert and A. Tsybakov. Fast learning rates for plug-in classifiers. *Ann. Statist.*, 35(2):608–633, 2007.
- [CCB16] A. Choromanska, K. Choromanski, and M. Bojarski. On the boosting ability of top-down decision tree learning algorithm for multiclass classification. preprint, 2016.
- [Cho70] C.K. Chow. On optimum error and reject trade-off. *IEEE Transactions on Information Theory*, 16:41–46, 1970.
- [dCDB09] J. del Coz, J. Díez, and A. Bahamonde. Learning nondeterministic classifiers. *J. Mach. Learn. Res.*, 10:2273–2293, 2009.
- [DH15] Ch. Denis and M. Hebiri. Consistency of plug-in confidence sets for classification in semi-supervised learning. preprint, 2015.
- [HW06] R. Herbei and M. Wegkamp. Classification with reject option. *Canad. J. Statist.*, 34(4):709–721, 2006.
- [Lei14] J. Lei. Classification with confidence. *Biometrika*, 101(4):755–769, 2014.
- [LRW13] J. Lei, J. Robins, and L. Wasserman. Distribution-free prediction sets. *J. Amer. Statist. Assoc.*, 108(501):278–287, 2013.
- [Tsy04] A. Tsybakov. Optimal aggregation of classifiers in statistical learning. *Ann. Statist.*, 32(1):135–166, 2004.
- [Vap98] V. Vapnik. *Statistical learning theory*. Adaptive and Learning Systems for Signal Processing, Communications, and Control. John Wiley & Sons Inc., New York, 1998. A Wiley-Interscience Publication.
- [vdV98] A. van der Vaart. *Asymptotic statistics*, volume 3 of *Cambridge Series in Statistical and Probabilistic Mathematics*. Cambridge University Press, Cambridge, 1998.
- [VGS99] V. Vovk, A. Gammerman, and C. Saunders. Machine-learning applications of algorithmic randomness. In *Proceedings of the 16th International Conference on Machine Learning*, pages 444–453. 1999.
- [VGS05] V. Vovk, A. Gammerman, and G. Shafer. *Algorithmic learning in a random world*. Springer, New York, 2005.
- [Vov02] V. Vovk. Asymptotic optimality of transductive confidence machine. In *Algorithmic learning theory*, volume 2533 of *Lecture Notes in Comput. Sci.*, pages 336–350. Springer, Berlin, 2002.
- [WLW04] T.-F. Wu, C.-J. Lin, and R. Weng. Probability estimates for multi-class classification by pairwise coupling. *J. Mach. Learn. Res.*, 5:975–1005, 2003/04.
- [WY11] M. Wegkamp and M. Yuan. Support vector machines with a reject option. *Bernoulli*, 17(4):1368–1385, 2011.
- [YW10] M. Yuan and M. Wegkamp. Classification methods with reject option based on convex risk minimization. *J. Mach. Learn. Res.*, 11:111–130, 2010.