



On Generalizing the C-Bound to the Multiclass and Multi-label Settings

François Laviolette, Emilie Morvant, Liva Ralaivola, Jean-Francis Roy

► To cite this version:

François Laviolette, Emilie Morvant, Liva Ralaivola, Jean-Francis Roy. On Generalizing the C-Bound to the Multiclass and Multi-label Settings. NIPS 2014 Workshop on Representation and Learning Methods for Complex Outputs, Dec 2014, Montréal, Canada. hal-01093560v2

HAL Id: hal-01093560

<https://hal.science/hal-01093560v2>

Submitted on 13 Jan 2015

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.

On Generalizing the C-Bound to the Multiclass and Multi-label Settings

François Laviolette

IFT-GLO, Université Laval
Québec (QC), Canada

francois.laviolette@ift.ulaval.ca

Emilie Morvant

Univ. de St-Etienne, France
LaHC, UMR CNRS 5516

emilie.morvant@univ-st-etienne.fr

Liva Ralaivola

Aix-Marseille Université
LIF, UMR CNRS 7279

liva.ralaivola@lif.univ-mrs.fr

Jean-Francis Roy

IFT-GLO, Université Laval
Québec (QC), Canada

jean-francis.roy@ift.ulaval.ca

Abstract

The \mathcal{C} -bound, introduced in Lacasse et al. [1], gives a tight upper bound on the risk of a binary majority vote classifier. In this work, we present a first step towards extending this work to more complex outputs, by providing generalizations of the \mathcal{C} -bound to the multiclass and multi-label settings.

1 Introduction

In binary classification, many state-of-the-art algorithms output prediction functions that can be seen as a majority vote of “simple” classifiers. Ensemble methods such as Bagging [2], Boosting [3] and Random Forests [4] are well-known examples of learning algorithms that output majority votes. Majority votes are also central in the Bayesian approach (see Gelman et al. [5] for an introductory text); in this setting, the majority vote is generally called the *Bayes Classifier*. It is also -interesting to point out that classifiers produced by kernel methods, such as the Support Vector Machine [6], can also be viewed as majority votes. Indeed, to classify an example x , the SVM classifier computes $\text{sign}\left(\sum_{i=1}^{|S|} \alpha_i y_i k(x_i, x)\right)$. Hence, as for standard binary majority votes, if the total weight of each $\alpha_i y_i k(x_i, x)$ that votes positive is larger than the total weight for the negative choice, the classifier will output a $+1$ label (and a -1 label in the opposite case).

Most bounds on majority votes take into account the margin of the majority vote on an example (x, y) , that is the difference between the total vote weight that has been given to the winning class minus the weight given to the alternative class. As an example, PAC-Bayesian bounds give bounds on majority votes classifiers by relating it to a stochastic classifier, called the *Gibbs* classifier which is, up to a linear transformation equivalent to the first statistical moment of the margin when (x, y) is drawn *i.i.d.* from a distribution [7]. Unfortunately, in most ensemble methods, the voters are weak and no majority vote can obtain high margins. Lacasse et al. [1] proposed a tighter relation between the risk of the majority vote that take into account both the first and the second moments of the margin: the \mathcal{C} -bound. This sheds a new light on the behavior of majority votes: it is not only how good are the voters but also how they are correlated in their voting. Namely, this has inspired a new learning algorithm named MinCq [7], whose performance is state-of-the-art. In this work, we generalize the \mathcal{C} -bound for multiclass and multi-label weighted majority votes as a first step towards the goal of designing learning algorithms for more complex outputs.

This paper is organized as follows. Section 2 recalls the \mathcal{C} -bound in binary classification. We generalize it to the multiclass and multi-label settings in Sections 3 and 4. We conclude in Section 5.

2 The \mathcal{C} -bound for Binary Classification

In this section, we recall the \mathcal{C} -bound [1, 7] in the binary classification setting.

Let $X \subseteq \mathbb{R}^d$ be the input space of dimension d , and let $Y = \{-1, +1\}$ be the output space. The learning sample $S = \{(x_i, y_i)\}_{i=1}^m$ is constituted by m examples drawn *i.i.d.* from a fixed but unknown distribution D over $X \times Y$. Let \mathcal{H} be a set of real-valued voters from X to Y . Given a prior distribution π on \mathcal{H} and given S , the goal of the PAC-Bayesian approach is to find the posterior distribution ρ on \mathcal{H} which minimizes the true risk of the ρ -weighted majority vote $B_\rho(\cdot)$ given by

$$\mathbf{R}_D(B_\rho) = \mathbf{E}_{(x,y) \sim D} \mathbf{I}(B_\rho(x) \neq y), \quad \text{where} \quad B_\rho(x) = \text{sign} \left[\mathbf{E}_{h \sim \rho} h(x) \right],$$

and where $\mathbf{I}(a) = 1$ if predicate a is true and 0 otherwise.

It is well-known that minimizing $\mathbf{R}_D(B_\rho)$ is NP-hard. To get around this problem, one solution is to make use of the \mathcal{C} -bound which is a tight bound over $\mathbf{R}_D(B_\rho)$. This bound is based on the notion of margin of $B_\rho(\cdot)$ defined as follows.

Definition 1 (the margin). *Let M_ρ^D be the random variable that, given an example (x, y) drawn according to D , outputs the margin of $B_\rho(\cdot)$ on that example, defined by $M_\rho(x, y) = y \mathbf{E}_{h \sim \rho} h(x)$.*

We then consider the first and second statistical moments of the random variable M_ρ^D , respectively given by $\mu_1(M_\rho^D) = \mathbf{E}_{(x,y) \sim D} M_\rho(x, y)$ and $\mu_2(M_\rho^D) = \mathbf{E}_{(x,y) \sim D} (M_\rho(x, y))^2$.

According to the definition of the margin, $B_\rho(\cdot)$ correctly classifies an example (x, y) when its margin is strictly positive, i.e. $\mathbf{R}_D(B_\rho) = \mathbf{Pr}_{(x,y) \sim D} (M_\rho(x, y) \leq 0)$. This equality makes it possible to prove the following theorem.

Theorem 1 (The \mathcal{C} -bound of Laviolette et al. [7]). *For every distribution ρ on a set of real-valued functions \mathcal{H} , and for every distribution D on $X \times Y$, if $\mu_1(M_\rho^D) > 0$, then we have*

$$\mathbf{R}_D(B_\rho) \leq 1 - \frac{(\mu_1(M_\rho^D))^2}{\mu_2(M_\rho^D)}.$$

Proof. The Cantelli-Chebyshev inequality states that any random variable Z and any $a > 0$, we have that $\mathbf{Pr}(Z \leq \mathbf{E}[Z] - a) \leq \frac{\text{Var } Z}{\text{Var } Z + a^2}$. We obtain the result by applying this inequality with $Z = M_\rho(x, y)$, and with $a = \mu_1(M_\rho^D)$, and by using the definition of the variance. \square

Note that the minimization of the empirical counterpart of the \mathcal{C} -bound is a natural solution for learning a distribution ρ that leads to a ρ -weighted majority vote $B_\rho(\cdot)$ with low error. This strategy is justified thanks to an elegant PAC-Bayesian generalization bound over the \mathcal{C} -bound, and have led to a simple learning algorithm called MinCq [7].

In the following, we generalize this important theoretical result in the PAC-Bayesian literature to the multiclass setting.

3 Generalizations of the \mathcal{C} -bound for Multiclass Classification

In this section, we stand in the multiclass classification setting where the input space is still $X \subseteq \mathbb{R}^d$, but the output space is $Y = \{1, \dots, Q\}$, with a finite number of classes $Q \geq 2$. Let \mathcal{H} be a set of multiclass voters from X to Y . We recall that given a prior distribution π over \mathcal{H} and given a sample S , *i.i.d.* from D , the PAC-Bayesian approach looks for the ρ distribution which minimizes the true risk of the majority vote $B_\rho(\cdot)$. In the multiclass classification setting, $B_\rho(\cdot)$ is defined by

$$B_\rho(x) = \underset{c \in Y}{\text{argmax}} \left[\mathbf{E}_{h \sim \rho} \mathbf{I}(h(x) = c) \right]. \quad (1)$$

As in binary classification, the risk $\mathbf{R}_D(B_\rho)$ of a ρ -weighted majority vote can be related to the notion of margin realized on an example (x, y) . However, in multiclass classification, such a notion can be expressed in a variety of manners. In the next section, we present three versions of multiclass margins that are equivalent in binary classification.

3.1 Margins in Multiclass Classification

We first make use of the multiclass margin proposed by Breiman [4] for the random forests, which can be seen as the usual notion of margin. Note that when $Y = \{-1, +1\}$, we recover the usual notion of binary margin of Definition 1.

Definition 2 (the multiclass margin). Let D be a distribution over $X \times Y$, let \mathcal{H} be a set of multiclass voters. Given a distribution ρ on \mathcal{H} , the margin of the majority vote $B_\rho(\cdot)$ realized on $(x, y) \sim P$ is

$$M_\rho(x, y) = \mathbf{E}_{h \sim \rho} \mathbf{I}(h(x) = y) - \max_{c \in Y, c \neq y} \left(\mathbf{E}_{h \sim \rho} \mathbf{I}(h(x) = c) \right).$$

Like in the binary classification framework presented in Section 2, the majority vote $B_\rho(\cdot)$ correctly classifies an example if its ρ -margin is strictly positive, i.e., $\mathbf{R}_D(B_\rho) = \Pr_{(x, y) \sim D} (M_\rho(x, y) \leq 0)$.

We then consider the relaxation proposed by Breiman [4] that is based on a notion of *strength* of the majority vote in regard to a class c .

Definition 3 (the strength). Let \mathcal{H} be a set of multiclass voters from X to Y and let ρ be a distribution on \mathcal{H} . Let $S_{\rho, c}^D$ be the random variable that, given an example (x, y) drawn according to a distribution D over $X \times Y$, outputs the strength of the majority vote $B_\rho(\cdot)$ on that example according to a class $c \in Y$, defined by $S_{\rho, c}(x, y) = \mathbf{E}_{h \sim \rho} \mathbf{I}(h(x) = y) - \mathbf{E}_{h \sim \rho} \mathbf{I}(h(x) = c)$.

From this definition, one can show that

$$\mathbf{R}_D(B_\rho) = \Pr_{(x, y) \sim D} (M_\rho(x, y) \leq 0) \leq \sum_{c=1}^Q \Pr_{(x, y) \sim D} (S_{\rho, c}(x, y) \leq 0) - 1, \quad (2)$$

where we have the equality in the binary classification setting. Lastly, we consider a relaxation of the notion of margin, that we call the ω -margin.

Definition 4 (the ω -margin). Let \mathcal{H} be a set of multiclass voters from X to Y , let ρ be a distribution on \mathcal{H} and let $\omega \leq 1$. Let $M_{\rho, \omega}^D$ be the random variable that, given an example $(x, y) \sim D$ over $X \times Y$, outputs the ω -margin of the majority vote $B_\rho(\cdot)$ on that example, defined by

$$M_{\rho, \omega}(x, y) = \mathbf{E}_{h \sim \rho} \mathbf{I}(h(x) = y) - 1/\omega. \quad (3)$$

This notion of margin can be seen as the difference between the weight given by the majority vote to the correct class y and a certain threshold $1/\omega$. In the case of the binary classification, we have that the sign of the ω -margin with $\omega = 2$ is the same than the sign of the binary margin. This observation comes from the fact that $\mathbf{E}_{h \sim \rho} \mathbf{I}(h(x) = y)$ is the proportion of voters that vote y . In the binary case, this proportion is $\leq \frac{1}{2}$ when the majority vote makes a mistake, and $> \frac{1}{2}$ otherwise. The following theorem relates the risk of $B_\rho(\cdot)$ and the ω -margin associated to ρ .

Theorem 2. Let $Q \geq 2$ be the number of classes. For every distribution D over $X \times Y$ and for every distribution ρ over a set of multiclass voters \mathcal{H} , we have

$$\Pr_{(x, y) \sim D} (M_{\rho, Q}(x, y) \leq 0) \leq \mathbf{R}_D(B_\rho) \leq \Pr_{(x, y) \sim D} (M_{\rho, 2}(x, y) \leq 0). \quad (4)$$

Proof. First, let us prove the left-hand side inequality. We have

$$\begin{aligned} \mathbf{R}_D(B_\rho) &= \Pr_{(x, y) \sim D} (M_\rho(x, y) \leq 0) = \Pr_{(x, y) \sim D} \left(\mathbf{E}_{h \sim \rho} \mathbf{I}(h(x) = y) \leq \max_{c \in Y, c \neq y} \mathbf{E}_{h \sim \rho} \mathbf{I}(h(x) = c) \right) \\ &\geq \Pr_{(x, y) \sim D} \left(\mathbf{E}_{h \sim \rho} \mathbf{I}(h(x) = y) \leq \max_{c \in Y, c \neq y} \mathbf{E}_{h \sim \rho} \mathbf{I}(h(x) = c) \right) \\ &= \Pr_{(x, y) \sim D} \left(\mathbf{E}_{h \sim \rho} \mathbf{I}(h(x) = y) \leq \frac{1}{Q-1} \left[1 - \mathbf{E}_{h \sim \rho} \mathbf{I}(h(x) = y) \right] \right) = \Pr_{(x, y) \sim D} (M_{\rho, Q}(x, y) \leq 0). \end{aligned}$$

The right-hand side inequality is easily verified by observing that the majority vote necessarily makes a correct prediction if the weight given to the correct class y is higher than $\frac{1}{2}$. \square

All the above-mentioned notions of margin are equivalent if we stand in the binary classification setting. However, they differ in the multiclass setting. The multiclass margin of Definition 2 is associated to the true decision function in multiclass classification, and is calculated considering all other classes. The strength of Definition 3 depends on the true class y of x and corresponds to a combination of binary margins (one class versus another class) for $c \neq y$. The ω -margin of Definition 4 also depends on the true class y of x , but does not consider the other classes. This measure is easier to manipulate, but implies a higher indecision region (see Theorem 2).

3.2 Generalizations of the \mathcal{C} -bound in the Multiclass Setting

The following bound is based on the definition of the multiclass margin in multiclass (Definition 2).

Theorem 3 (the multiclass \mathcal{C} -bound). *For every distribution ρ on a set of multiclass voters \mathcal{H} , and for every distribution D on $X \times Y$, such that $\mu_1(M_\rho^D) > 0$, we have*

$$\mathbf{R}_D(B_\rho) \leq 1 - \frac{(\mu_1(M_\rho^D))^2}{\mu_2(M_\rho^D)}.$$

Proof. The proof is the same than the one of the binary \mathcal{C} -bound (see Theorem 1), by considering the multiclass majority-vote of Equation (1) and the multiclass margin of Definition 2. \square

This bound offers an accurate relation between the risk of the majority vote and the margin. However, the max term in the definition of the multiclass margin makes the derivation of an algorithm to minimize this bound much harder than in binary classification.

Thanks to the definition of the strength of Definition 3 and according to the proof process of the \mathcal{C} -bound, we obtain the following relation.

Theorem 4. *For every distribution ρ on a set of multiclass voters \mathcal{H} , and for every distribution D over $X \times Y$, such that $\forall c \in Y, \mu_1(S_{\rho,c}^D) > 0$, we have*

$$\mathbf{R}_D(B_\rho) \leq \sum_{c=1}^Q \mathbf{Pr}_{(x,y) \sim D} (S_{\rho,c}(x,y) \leq 0) - 1 = (Q-1) - \sum_{c=1}^Q \frac{(\mu_1(S_{\rho,c}^D))^2}{\mu_2(S_{\rho,c}^D)},$$

Proof. The result is obtained by using Inequality (2) in the proof of the \mathcal{C} -bound. \square

This result can be seen as a sum of \mathcal{C} -bounds for every class. A practical drawbacks of this bound in order to construct a minimization algorithm is that we have to minimize a sum of ratios. Finally, the \mathcal{C} -bound obtained by using the ω -margin is given by the following theorem.

Theorem 5. *For every distribution ρ on a set of multiclass voters \mathcal{H} , for every $\omega \geq 1$, and for every distribution D on $X \times Y$, if $\mu_1(M_{\rho,\omega}^D) > 0$, we have*

$$\mathbf{E}_{(x,y) \sim D} \mathbf{I}(M_{\rho,\omega}(x,y) \leq 0) \leq 1 - \frac{(\mu_1(M_{\rho,\omega}^D))^2}{\mu_2(M_{\rho,\omega}^D)}.$$

Proof. The result is obtained with the same proof process than the \mathcal{C} -bound, by replacing the use of the random variable M_ρ^D by $M_{\rho,\omega}^D$. \square

The ω -margin being linear, we are now able to build a bound minimization algorithm as in Laviolette et al. [7] for the multiclass classification setting.

4 Extending the ω -margin to the Multi-label Setting

In this section, we will extend the ω -margin with $\omega = 2$ to the more general *multi-label* classification setting. Doing so, we will be able to upper bound the risk of the multi-label majority vote classifier. We stand in the multi-label classification setting where the input space is still $X \subseteq \mathbb{R}^d$, the space of possible labels is $Y = \{1, \dots, Q\}$ with a finite number of classes $Q \geq 2$, but we consider the output space $\bar{Y} = \{0, 1\}^Q$ that contains vectors \bar{y} of length Q where the i^{th} element is 1 if example i is among the labels associated to the example x , and 0 otherwise. We consider a set $\bar{\mathcal{H}}$ of *multi-label voters* $\bar{h} : X \mapsto \bar{Y}$. As usual in structured output prediction, given a distribution ρ over $\bar{\mathcal{H}}$, the multi-label majority vote classifier \bar{B}_ρ chooses the label $\bar{c} \in \bar{Y}$ that has the lowest squared Euclidean distance with the ρ -weighted cumulative confidence,

$$\bar{B}_\rho(x) = \operatorname{argmin}_{\bar{c} \in \bar{Y}} \left\| \bar{c} - \mathbf{E}_{\bar{h} \sim \rho} \bar{h}(x) \right\|^2 = \operatorname{argmax}_{\bar{c} \in \bar{Y}} \left[\bar{c} \cdot \left(\mathbf{E}_{\bar{h} \sim \rho} \bar{h}(x) - \frac{1}{2} \mathbf{1} \right) \right],$$

where $\mathbf{1}$ is a vector of length Q containing ones. The multi-label margin is given by Definition 5.




Figure 1: Graphical representation of label \bar{y} , hyperplane $P_{\bar{y}}$ and vector $\mathbf{E}_{\bar{h} \sim \rho} \bar{h}(x)$.

Definition 5 (the multi-label margin). Let D be a distribution over $X \times \bar{Y}$, let $\bar{\mathcal{H}}$ be a set of multi-label voters. Given a distribution ρ on $\bar{\mathcal{H}}$, the margin of the majority vote $\bar{B}_\rho(\cdot)$ on (x, \bar{y}) is

$$\bar{M}_\rho(x, \bar{y}) = \left(\mathbf{E}_{\bar{h} \sim \rho} \bar{h}(x) - \frac{1}{2} \mathbf{1} \right) \cdot \left(\bar{y} - \frac{1}{2} \mathbf{1} \right) - \max_{\bar{c} \in \bar{Y}, \bar{c} \neq \bar{y}} \left(\mathbf{E}_{\bar{h} \sim \rho} \bar{h}(x) - \frac{1}{2} \mathbf{1} \right) \cdot \left(\bar{c} - \frac{1}{2} \mathbf{1} \right).$$

As we did in the multiclass setting with the margin of Definition 2, we can upper bound the risk of the multi-label majority vote classifier by developing a \mathcal{C} -bound using the margin of Definition 5. However, as this margin also depends on a max term, the derivation of a learning algorithm minimizing the resulting \mathcal{C} -bound remains hard. To overcome this, we generalize $M_{\rho,2}$, the ω -margin with $\omega = 2$ of Definition 4, to the multi-label setting, as follows.

Definition 6. Let D be a distribution over $X \times \bar{Y}$, let $\bar{\mathcal{H}}$ be a set of multi-label voters. Given a distribution ρ on $\bar{\mathcal{H}}$, the 2-margin of the majority vote $\bar{B}_\rho(\cdot)$ on (x, \bar{y}) is

$$\begin{aligned} \bar{M}_{\rho,2}(x, \bar{y}) &= \left(\left(\mathbf{E}_{\bar{h} \sim \rho} \bar{h}(x) - \frac{1}{2} \mathbf{1} \right) - \left(\bar{y}_{i \rightarrow 1/2} - \frac{1}{2} \mathbf{1} \right) \right) \cdot \left(\bar{y} - \frac{1}{2} \mathbf{1} \right) \\ &= \bar{y} \cdot \left(\mathbf{E}_{\bar{h} \sim \rho} \bar{h}(x) - \frac{1}{2} \mathbf{1} \right) - \mathbf{E}_{\bar{h} \sim \rho} \bar{h}(x) \cdot \frac{1}{2} \mathbf{1} - \frac{1}{4}, \end{aligned}$$

where $i \in \{1, \dots, Q\}$ and $\bar{y}_{i \rightarrow 1/2}$ is obtained from \bar{y} by replacing its i^{th} coordinate by $1/2$.

The second equality of the definition is obtained by straightforward calculation. Now, let $P_{\bar{y}}$ be the only hyperplane on which lies all the points of the form $\bar{y}_{i \rightarrow 1/2}$ for $i = 1, \dots, Q$. Since this hyperplane has normal $\left(\bar{y} - \frac{1}{2} \mathbf{1} \right)$, it follows from basic linear algebra that if $\bar{M}_{\rho,2} > 0$, then vectors $\mathbf{E}_{\bar{h} \sim \rho} \bar{h}(x)$ and \bar{y} will be on the same side of $P_{\bar{y}}$. It is also easy to see that in this case, we have $\bar{B}_\rho(x) = \bar{y}$. Figure 1 shows an example in the case where $Q = 2$. Thus, we have that $\mathbf{R}_D(\bar{B}_\rho) \leq \mathbf{Pr}_{(x, \bar{y}) \sim D} (\bar{M}_{\rho,2}(x, \bar{y}) \leq 0)$, and following the same arguments as in Theorem 5, one can derive the following multi-label \mathcal{C} -bound.

Theorem 6. For every distribution ρ on a set of multi-label voters $\bar{\mathcal{H}}$ and for every distribution D on $X \times \bar{Y}$, if $\mu_1(\bar{M}_{\rho,2}(x, \bar{y})) > 0$, we have

$$\mathbf{R}_D(\bar{B}_\rho) \leq \mathbf{E}_{(x, \bar{y}) \sim D} \mathbf{I}(\bar{M}_{\rho,2}(x, \bar{y}) \leq 0) \leq 1 - \frac{\left(\mu_1(\bar{M}_{\rho,2}^D) \right)^2}{\mu_2(\bar{M}_{\rho,2}^D)}.$$

5 Conclusion and Outlooks

In this paper, we extend an important theoretical result in the PAC-Bayesian literature to the multiclass and multi-label settings. Concretely, we prove three multiclass versions and one multi-label version of the \mathcal{C} -bound, a bound over the risk of the majority vote, based on generalizations of the notion of margin for multiclass and multi-label classification. These results open the way to extending the theory to more complex outputs and developing new algorithms for multiclass and multi-label classification with PAC-Bayesian generalization guarantees.

References

- [1] A. Lacasse, F. Laviolette, M. Marchand, P. Germain, and N. Usunier. PAC-bayes bounds for the risk of the majority vote and the variance of the Gibbs classifier. In *Proceedings of annual conference on Neural Information Processing Systems*, 2007.
- [2] Leo Breiman. Bagging predictors. *Machine Learning*, 24(2):123–140, 1996.
- [3] R. E. Schapire and Y. Singer. Improved boosting algorithms using confidence-rated predictions. In *Machine Learning Journal*, pages 80–91, 1999.
- [4] Leo Breiman. Random Forests. *Machine Learning*, 45(1):5–32, October 2001.
- [5] A. Gelman, J.B. Carlin, H.S. Stern, and D.B. Rubin. *Bayesian Data Analysis*. Chapman & Hall/CRC, 2004.
- [6] Corinna Cortes and Vladimir Vapnik. Support-vector networks. *Machine Learning*, 20(3):273–297, 1995.
- [7] F. Laviolette, M. Marchand, and J.-F. Roy. From PAC-Bayes Bounds to Quadratic Programs for Majority Votes. In *Proceedings of International Conference on Machine Learning*, June 2011.