
On the Consistency of Max-Margin Losses

Alex Nowak-Vila
ENS-INRIA-PSL
Paris, France

Alessandro Rudi
ENS-INRIA-PSL
Paris, France

Francis Bach
ENS-INRIA-PSL
Paris, France

Abstract

The foundational concept of Max-Margin in machine learning is ill-posed for output spaces with more than two labels such as in structured prediction. In this paper, we show that the Max-Margin loss can only be consistent to the classification task under highly restrictive assumptions on the discrete loss measuring the error between outputs. These conditions are satisfied by distances defined in tree graphs, for which we prove consistency, thus being the first losses shown to be consistent for Max-Margin beyond the binary setting. We finally address these limitations by correcting the concept of Max-Margin and introducing the Restricted-Max-Margin, where the maximization of the loss-augmented scores is maintained, but performed over a subset of the original domain. The resulting loss is also a generalization of the binary support vector machine and it is consistent under milder conditions on the discrete loss.

1 INTRODUCTION

One of the first binary classification methods learned in a machine learning course is the support vector machine (SVM) (Boser et al., 1992; Cortes and Vapnik, 1995) and it is introduced using the principle of maximum margin: assuming the data are linearly separable, the classification hyperplane must maximize the separation to the observed examples. Having this intuition in mind, the same principle has been used to extend this notion to larger output spaces \mathcal{Y} , such as multi-class classification (Crammer and

Singer, 2001) and structured prediction (Taskar et al., 2004; Tsochantaridis et al., 2005), where the separation to the observed examples is controlled by a discrete loss $L(y, y')$ measuring the error between outputs y and y' . The resulting method generalizes the binary SVM and corresponds to minimizing the so-called Max-Margin loss

$$S_M(v, y) = \max_{y' \in \mathcal{Y}} L(y, y') + v_{y'} - v_y, \quad (1)$$

where $v \in \mathbb{R}^{|\mathcal{Y}|}$ is a vector with coordinate v_y encoding the score for output y . Unfortunately, this method may not be *consistent*, i.e., minimizing the Max-Margin loss (1) may not lead to a minimization of the discrete loss L of interest. In particular, it is known that the Max-Margin loss is only consistent for the 0-1 loss under the dominant label condition, i.e., when for every input there exists an output element with probability larger than 1/2 (Liu, 2007), which is always satisfied in the binary case. However, far less is known for other tasks. Indeed, the Max-Margin loss is widely used for structured output spaces where the discrete loss L defining the task is different than the 0-1 loss, under the name of Structural SVM (SSVM) (Taskar et al., 2005; Caetano et al., 2009; Smith, 2011) or Max-Margin Markov Networks (M³N) (Taskar et al., 2005). In this general setting, the following questions remain unanswered:

- (i) *Does there exist a necessary condition on L for consistency to hold? Does it exist a space of losses for which consistency holds? Can we generalize the consistency result under the dominant label assumption beyond the 0-1 loss?*
- (ii) *Can we correct the Max-Margin loss to make it consistent by maintaining the additive and maximization structure of the Max-Margin loss?*

We answer these questions in this paper. In particular, we make the following contributions:

- We prove that the Max-Margin loss can only be consistent under a restrictive necessary condition

on the structure of the loss L , indeed, the loss L has to be a distance and satisfy the triangle inequality as an *equality* for several groups of outputs. As a positive result, we show that a distance defined in a tree graph, such as the absolute deviation loss used in ordinal regression, satisfies this condition and it is consistent, thus providing the first set of losses for which consistency holds beyond the binary setting. We also extend the existing *partial* consistency result of the 0-1 loss by extending the result under the dominant label condition to all losses that are distances.

- As a secondary contribution, we introduce the Restricted-Max-Margin loss, where the maximization of the loss-augmented scores defining the Max-Margin loss is restricted to a subset of the simplex. The resulting loss also generalizes the binary SVM and it is consistent under milder assumptions on L . Moreover, we show the connections between these losses and the Max-Min-Margin loss (Fathony et al., 2016; Duchi et al., 2018; Nowak-Vila et al., 2020), where consistency always holds independently of the discrete loss L .

2 MAX-MARGIN AND MAIN RESULTS

In this section we introduce the concept of Max-Margin learning and its consistency from its origins in binary classification to the structured output setting. This is followed by the presentation of the main results of this paper and its implications are discussed.

2.1 Max-Margin Learning

Binary output. Let $(x_1, y_1), \dots, (x_n, y_n)$ be n examples of input-output pairs sampled from an unknown distribution ρ defined in $\mathcal{X} \times \mathcal{Y}$. Let us first assume that the input space \mathcal{X} is a vector space and $\mathcal{Y} = \{-1, 1\}$ represents binary labels. The goal is to construct a binary-valued function $f : \mathcal{X} \rightarrow \mathcal{Y}$ minimizing the expected classification error

$$\mathcal{E}(f) = \mathbb{E}_{(x,y) \sim \rho} L(f(x), y), \quad (2)$$

where $L(y, y') = 1(y \neq y')$ is the binary 0-1 loss. The concept of max-margin was initially defined in this setting to construct a predictor of the form $\text{sign}(g(x))$ where $g(x) = w^\top x + b$ is an affine function defining a hyperplane with maximum separation to the examples assuming linearly separable data (Boser et al., 1992). In this setting, an example x_i is correctly classified if $(w^\top x_i + b_i)y_i > 0$ and misclassified otherwise. The max-margin hyperplane can be found by minimizing $\|w\|_2^2$ under the constraint $(w^\top x_i + b)y_i \geq 1$ for all

n examples. When the data are not linearly separable, some examples are allowed to be misclassified by introducing some non-negative *slack variables* ξ_i and solving the optimization problem known as the *support vector machine (SVM)* (Cortes and Vapnik, 1995):

$$\begin{cases} \min_{w,b,\xi} & \frac{1}{n} \sum_{i=1}^n \xi_i + \frac{\lambda}{2} \|w\|_2^2 \\ \text{s.t.} & (w^\top x_i + b)y_i \geq 1 - \xi_i, \quad \xi_i \geq 0, \quad \forall i \in [n], \end{cases}$$

where $\lambda > 0$ is a parameter used to balance the first term with the second. We can re-write the constraints as $\xi_i \geq 1 - y_i g(x_i)$ for non-negative ξ_i 's and extend the affine hypothesis space to a generic functional space \mathcal{G} with associated norm $\|\cdot\|_{\mathcal{G}}$ to allow for non-linear predictors, such as reproducing kernel Hilbert spaces (RKHS) (Aronszajn, 1950). Then, the problem above can be written as a convex regularized empirical risk minimization (ERM) (Vapnik, 1992) problem

$$\min_{g \in \mathcal{G}} \frac{1}{n} \sum_{i=1}^n S_M(g(x_i), y_i) + \frac{\lambda}{2} \|g\|_{\mathcal{G}}^2, \quad (3)$$

where $S_M(v, y) = \max(1 - yv, 0)$ is the binary Max-Margin loss (also called SVM loss), and now λ can be interpreted as the regularization parameter. An important property of the classification method is that the estimated predictor solving (3) over all measurable functions converges to the predictor f^* minimizing the expected classification error (2) in the infinite data regime ($n \rightarrow \infty$ and $\lambda \rightarrow 0$) (Vapnik, 2013). More concretely, the minimizer g^* of the expected risk $\mathbb{E}_{(x,y) \sim \rho} S_M(g(x), y)$ must satisfy $f^* = \text{sign}(g^*)$. This property is called *Fisher consistency* (Bartlett et al., 2006) (or simply consistency) and can be studied in terms of the conditional expectation $q(x) := \rho(1|x)$, as $f^*(x)$ and $g^*(x)$ can be characterized in terms of this quantity¹. Note that in the rest of the paper we will drop the dependence in x from the function q : a statement $P(q)$ for all $q \in [0, 1]$ must then be read as $P(q(x))$ for all $x \in \mathcal{X}$. Let $v_M^*(q) \subseteq \mathbb{R}$ and $y^*(q) \subseteq \mathcal{Y}$ be the minimizers of the conditional risks $\mathbb{E}_{y' \sim q} S_M(v, y')$ and $\mathbb{E}_{y' \sim q} L(y, y')$, respectively. Then, Fisher consistency is equivalent to say that if $v \in v_M^*(q)$, then $\text{sign}(v) \in y^*(q)$ for all $q \in [0, 1]$ (Devroye et al., 2013).

¹This is because f^* and g^* are minimizers over all measurable functions of an expectation over \mathcal{X} .

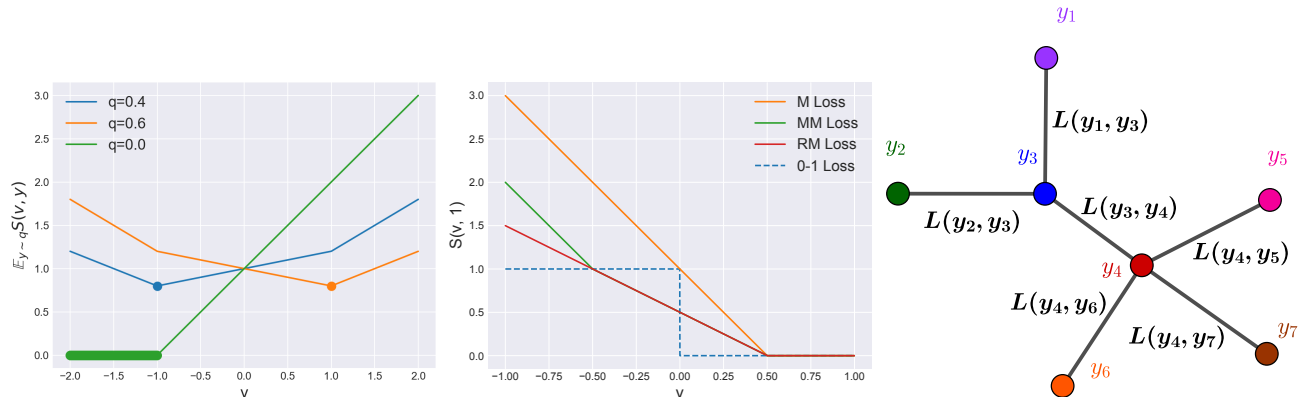


Figure 1: **Left:** Plots of the conditional risks of the Max-Margin loss for $q = 0.4$, $q = 0.6$ and $q = 0$, respectively. The conditional risk in the set of minimizers $v_M^*(q)$ is plotted with a thick point / line. **Middle:** Plot of $S_M(v, 1)$, $S_{MM}(v, 1)$ and $S_{RM}(v, 1)$ in the binary setting with $v = v_1 = -v_{-1}$. In this case, $S_M = 2S_{RM}$ (so both losses generalize the binary SVM up to a factor of 2) and S_M is the only one upper-bounding the 0-1 loss. Moreover, the three losses are consistent with the 0-1 loss. **Right:** Distance defined in a tree: the distance/loss between two nodes is the sum of the distances between adjacent nodes of the path between them. For every triplet of outputs $y, y', y'' \in \mathcal{Y}$, either they are aligned in a path, or there exists $z \in \mathcal{Y}$ belonging to the shortest path between all pairs.

This property is satisfied as (see also left Fig. 1):

$$y^*(q) = \begin{cases} \{1\} & q \in (1/2, 1] \\ \{-1, 1\} & q = 1/2 \\ \{-1\} & q \in [0, 1/2), \end{cases} \quad v_M^*(q) = \begin{cases} [1, \infty) & q = 1 \\ \{1\} & q \in (1/2, 1) \\ [-1, 1] & q = 1/2 \\ \{-1\} & q \in (0, 1/2) \\ (-\infty, -1] & q = 0. \end{cases}$$

Structured prediction. In the structured prediction setting, we have $k = |\mathcal{Y}|$ possible outputs and the goal is to estimate a discrete-valued function f minimizing (2) where now $L : \mathcal{Y} \times \mathcal{Y} \rightarrow \mathbb{R}$ is a generic non-negative discrete loss function between output pairs defining the task at hand. We construct predictors of the form $\arg \max_{y \in \mathcal{Y}} g_y(x)$, where $g : \mathcal{X} \rightarrow \mathbb{R}^k$ is a vector-valued function assigning scores to each of the k possible outputs. The maximum margin principle from binary classification is generalized as follows. For every example (x_i, y_i) , the method minimizes the squared norm $\|g\|_2^2$ under the constraints

$$g_{y_i}(x_i) \geq \underbrace{L(y, y_i) + g_y(x_i)}_{\text{loss-augmented scores}}, \quad (4)$$

for all possible outputs y . By writing the above constraint as $g_{y_i}(x_i) - g_y(x_i) \geq L(y, y_i)$, we observe that this generalizes the condition $y_i g(x_i) \geq 1$ from binary classification when $g = g_1 = -g_{-1}$ so that the argmax corresponds to the sign and L is the binary 0-1 loss. As in the binary case, introducing slack variables

and turning it into a regularized ERM problem of the form (3) we obtain the Max-Margin loss $S_M(v, y) = \max_{y' \in \mathcal{Y}} L(y, y') + v_{y'} - v_y$, which is constructed as a maximization of the loss-augmented scores defined in (4). To ease notation, the dependence of S on the loss L is deduced from the context. The Max-Margin loss (1) is known as the Crammer-Singer SVM (Crammer and Singer, 2001) when L is the 0-1 loss, and it is also widely used in structured prediction settings with exponentially large output spaces under the name of Structural SVM (Joachims, 2006) or Max-Margin Markov Networks (M³N) (Taskar et al., 2004) by using losses between structured outputs such as sequences, permutations, graphs, etc (BakIr et al., 2007). An interesting property of this loss is that it upper-bounds the discrete loss as $L(\arg \max_{y' \in \mathcal{Y}} v_{y'}, y) \leq S_M(v, y)$, for all $v \in \mathbb{R}^k$ and $y \in \mathcal{Y}$, which can guide us to think that minimizing $\mathbb{E}_{(x, y) \sim \rho} S_M(g(x), y)$ leads to minimizing $\mathbb{E}_{(x, y) \sim \rho} L(\arg \max_{y' \in \mathcal{Y}} g_{y'}(x), y)$. Unfortunately this intuition is misleading, as this bound is in general far from tight. Analogously to the binary case, let $q : \mathcal{X} \rightarrow \Delta$ be the conditional distribution where Δ is the simplex over \mathcal{Y} (we again drop the dependence on x , since each statement must be read as holding for every $x \in \mathcal{X}$). Moreover, we define for every q the set of minimizers of the conditional risks as

$$y^*(q) = \arg \min_{y \in \mathcal{Y}} L_y^\top q \subseteq \mathcal{Y}, \\ v_M^*(q) = \arg \min_{v \in \mathbb{R}^k} S_M(v)^\top q \subseteq \mathbb{R}^k,$$

where $L_y = (L(y, y'))_{y' \in \mathcal{Y}} \in \mathbb{R}^k$ is the y -th row of the loss matrix and $S_M(v) = (S_M(v, y'))_{y' \in \mathcal{Y}} \in \mathbb{R}^k$. We

say that S_M is Fisher consistent to L if for all $q \in \Delta$

$$v \in v_M^*(q) \implies \arg \max_{y \in \mathcal{Y}} v_y \in y^*(q). \quad (5)$$

Related works on consistency of Max-Margin.

The Max-Margin loss is only consistent to the 0-1 loss under the dominant label assumption $\max_{y \in \mathcal{Y}} q_y \geq 1/2$ (Liu, 2007). Ramaswamy et al. (2018) show that it is consistent to the ‘‘abstain’’ loss, but in this case the loss appearing in the definition (1) is not the same as the classification loss L . McAllester (2007) studies consistency of non-convex versions of max-margin methods on linear hypothesis spaces. There exist several generalizations of the binary SVM to larger output spaces other than Max-Margin (Dogan et al., 2016) such as Weston-Watkins (WW-SVM) (Weston and Watkins, 1999), Lee-Lin-Wahba (LLW-SVM) (Lee et al., 2004), Simplex-Coding (SC-SVM) (Mroueh et al., 2012), with the last two being consistent and defined as sums. However, the only loss with a max-structure is (1), which makes it computationally feasible to work in structured spaces of exponential size such as sequences or permutations. The Max-Min-Margin loss (Fathony et al., 2016; Duchi et al., 2018; Nowak-Vila et al., 2020) (defined below in Eq. (8)) is always consistent, it has a max-min structure and can be used in structured prediction settings. However, it does not correspond to the SVM in the binary setting, so it cannot be considered a generalization of the binary SVM.

2.2 Main Results

We assume that L is symmetric and that $L(y, y') = 0$ if and only if $y = y'$. Symmetry of L is assumed for the sake of exposition, but it is only required for the results on Max-Margin.

Main Results on Max-Margin. The following Thm. 2.1 is our main negative result.

Theorem 2.1 (Necessary condition for consistency S_M). *Let $k = |\mathcal{Y}| > 2$. If the Max-Margin loss is consistent to L , then L is a distance and for every three outputs $y_1, y_2, y_3 \in \mathcal{Y}$, there exists $z \in \mathcal{Y}$ for which the following three identities hold:*

$$\begin{aligned} L(y_1, y_2) &= L(y_1, z) + L(z, y_2), \\ L(y_1, y_3) &= L(y_1, z) + L(z, y_3), \\ L(y_2, y_3) &= L(y_2, z) + L(z, y_3). \end{aligned}$$

If $z = y_2$ in Thm. 2.1, then the only informative condition is $L(y_1, y_3) = L(y_1, y_2) + L(y_2, y_3)$ as L is assumed to be symmetric, which means that the outputs y_1, y_2, y_3 are ‘aligned’ in the output space (analogously

for $z = y_2, y_3$). On the other hand, if $z \neq y_1, y_2, y_3$, then the three equations are informative and all distances between the pairs can be decomposed into distances to z . The following discrete losses do not satisfy the above necessary condition (see Section 2 of Appendix):

- Losses which are not distances (such as the squared discrete loss $(y - y')^2$).
- Losses with full rank loss matrix with existing $q \in \text{int}(\Delta)$ for which all outputs are optimal, i.e., $y^*(q) = \mathcal{Y}$ (such as the 0-1 loss).
- Hamming losses $L(y, y') = \frac{1}{M} \sum_{m=1}^M L_m(y_m, y'_m)$ with $y, y' \in \prod_{m=1}^M \mathcal{Y}_m$ where L_m does not satisfy the necessary conditions for some $m = 1, \dots, M$.
- Hamming loss on permutations $L(\sigma, \sigma') = \frac{1}{M} \sum_{m=1}^M 1(\sigma(m) \neq \sigma'(m))$ with σ, σ' permutations of size M , used for graph matching (Pettersson et al., 2009; Caetano et al., 2009).

It is an open question whether the necessary condition of Thm. 2.1 is also sufficient. The following Thm. 2.2 shows that distances defined in a tree, which always satisfy this condition (see right Fig. 1), are indeed consistent.

Theorem 2.2 (Sufficient condition for consistency S_M). *If L is a distance defined in a tree, then the Max-Margin loss is consistent to L .*

An important example of these losses is the *absolute deviation* loss used in ordinal regression,

$$L(y, y') = |\gamma_y - \gamma_{y'}|, \quad \gamma \in \mathbb{R}^k,$$

for which the associated tree is a chain. Note that these losses are not the only ones satisfying the necessary condition given by Thm. 2.1. Indeed, the Hamming loss with $M = 2$, $\mathcal{Y}_1 = \mathcal{Y}_2 = \{-1, 1\}$ and L_1, L_2 the 0-1 loss is *not* a distance in a tree, satisfies the necessary condition, and consistency can be proven to hold (see Section 2 of Appendix). The following Prop. 2.3 gives a much milder sufficient condition to ensure *partial* consistency under the dominant label assumption, thus generalizing the well-known results from Liu (2007).

Proposition 2.3 (Sufficient condition for *partial* consistency S_M). *If L is a distance, then the Max-Margin loss is consistent to L under the dominant label assumption, i.e., $\max_{y \in \mathcal{Y}} q_y \geq 1/2$.*

In other words, if the learning task is defined by a distance and it is close to deterministic, then the Max-Margin loss is consistent to the task.




Figure 2: From left to right: plots of $S_M(v, 1)$, $S_{MM}(v, 1)$ and $S_{RM}(v, 1)$ in the three-label setting with $v^\top \mathbf{1} = 0$ for the 0-1 loss (note that $S(v + c\mathbf{1}, y) = S(v, y)$ for all three losses). The Max-Min-Margin and the Restrictive-Max-Margin loss coincide in the bottom-left region, but the Max-Min-Margin loss has three extra activated faces in the top-right region of the plot which are in general unnecessary for consistency, while the Restricted-Max-Margin uses just the necessary ones. The Max-Margin in the left plot uses just three faces, being insufficient for consistency.

Beyond Max-Margin. To overcome the limitations imposed by the maximum margin, but retaining the maximization structure of the loss, we propose a novel generalization of the binary SVM to structured prediction by restricting the maximization of the loss-augmented scores in (1). First, note that the Max-Margin loss can be written as a maximization over the simplex Δ over \mathcal{Y} as

$$S_M(v, y) = \max_{q \in \Delta} L_y^\top q + v^\top q - v_y. \quad (6)$$

We restrict the maximization to the so-called *prediction set* $\Delta(y) = \{q \in \Delta \mid y \in y^*(q)\}$, defined as the set of probabilities for which y is optimal. In binary classification the sets are $\Delta(-1) = [0, 1/2]$ and $\Delta(1) = [1/2, 1]$. The resulting *Restricted-Max-Margin* loss reads

$$S_{RM}(v, y) = \max_{q \in \Delta(y)} L_y^\top q + v^\top q - v_y. \quad (7)$$

This loss satisfies $2S_{RM} = S_M$ in the binary setting (see middle Fig. 1), thus, it corresponds to the binary SVM up to a scaling with a factor of two. The following Thm. 2.4 states that consistency of S_M implies consistency of S_{RM} and provides a sufficient condition for consistency of S_{RM} .

Theorem 2.4 (Sufficient condition for consistency S_{RM}). *The Restricted-Max-Margin loss is consistent to L whenever the Max-Margin is consistent. Moreover, if L satisfies $q_y > 0$ for every y optimal for $q \in \Delta$, i.e., $q \in \Delta(y)$, then the Restricted-Max-Margin is also consistent to L .*

In other words, if the output y is optimal for q , then the probability of this label has to be strictly greater

than zero $q_y > 0$. The 0-1 loss, which does not satisfy the necessary condition of Thm. 2.1, satisfies the sufficient condition for the Restricted-Max-Margin, as $\min_{q \in \Delta(y)} q_y = 1/k$ for all $y \in \mathcal{Y}$. However, there are still losses for which (7) is not consistent to, such as the squared discrete loss $(z - y)^2$ (see right of Fig. 3). The remaining inconsistencies can be resolved by going beyond the maximization structure into a max-min structure. The resulting loss is the so-called *Max-Min-Margin loss* (Fathony et al., 2016; Duchi et al., 2018; Nowak-Vila et al., 2020) defined as

$$S_{MM}(v, y) = \max_{q \in \Delta} \min_{z \in \mathcal{Y}} L_z^\top q + v^\top q - v_y. \quad (8)$$

It is known (Nowak-Vila et al., 2020) that the loss (8) is *always* consistent to L . As shown in Fig. 1 (middle), this loss does not correspond to the SVM in the binary setting because it has two symmetric kinks instead of one. See also Fig. 2 to compare the shape of the different losses for $k = 3$. Hence, while the structure of the loss gets more computationally involved from the Max-Margin loss (6) to the Max-Min-Margin loss (8), passing by the Restricted-Max-Margin loss (7), the consistency properties of these losses improve from one to the next.

3 BACKGROUND AND PRELIMINARY RESULTS

3.1 Background on Polyhedral Losses

Fisher consistency. Let us now consider a generic loss $S : \mathbb{R}^k \times \mathcal{Y} \rightarrow \mathbb{R}$ and let's generalize the argmax computing the prediction from the scores in the previ-

ous section to a generic *decoding* function $d : \mathbb{R}^k \rightarrow \mathcal{Y}$. The set of minimizers $v^*(q) \subseteq \mathbb{R}^k$ of the conditional risk $S(v)^\top q$ is also defined as before. We say that S is *Fisher consistent* to L (Tewari and Bartlett, 2007) under the decoding $d : \mathbb{R}^k \rightarrow \mathcal{Y}$ if

$$v \in v^*(q) \implies d(v) \in y^*(q), \quad (9)$$

for all $q \in \Delta$. If the decoding is not specified, it means that there exists a decoding satisfying this property. An important quantity throughout the paper is the *Bayes risk*, a *concave* function defined as the minimum of the conditional expected loss respectively for L and S :

$$H_L(q) = \min_{y \in \mathcal{Y}} L_y^\top q, \quad H_S(q) = \min_{v \in \mathbb{R}^k} S(v)^\top q.$$

Embedding of discrete losses. Fisher consistency in Eq. (9) states that every minimizer $v \in v^*(q)$ can be assigned to a solution of the discrete task using the decoding d . For the analysis of this paper, it will be useful to work using the concept of embeddability between losses (Finocchiaro et al., 2019), a stronger notion than Fisher consistency.

Definition 3.1 (Embeddability). *S embeds L if there exists an embedding $\psi : \mathcal{Y} \rightarrow \mathbb{R}^k$ such that: (i) $y \in y^*(q) \iff \psi(y) \in v^*(q)$, $\forall q \in \Delta$, and (ii) $S(\psi(y)) = L_y$, $\forall y \in \mathcal{Y}$.*

Condition (i) states that every solution of the discrete problem corresponds to a solution of the problem in S and vice versa. In particular, this rules out many smooth plug-in classifiers such as the squared loss or logistic regression, because they predict the vector of probabilities q which cannot be recovered from the discrete predictor $y^*(q)$. It is known (Finocchiaro et al., 2019) that the existence of an embedding ψ satisfying (i) implies the existence of a decoding d satisfying Eq. (9), so it is already a sufficient condition for Fisher consistency. Note that both Eq. (9) and condition (i) are assumptions on the predictors y^* and v^* , but there exist several losses L with the same set of minimizers $y^*(q)$ of the conditional risk $L_y^\top q$. The same can be said for the objects v^* and S . Condition (ii) restricts the relationship between pairs of losses by assuming that L can be recovered from S using the embedding ψ . The following Prop. 3.2 shows that S embedding L is equivalent to having the same Bayes risks.

Proposition 3.2 (Finocchiaro et al. (2019)). *S embeds L if and only if $H_L = H_S$.*

Moreover, it is known that any discrete loss L is embedded by at least one loss S (Theorem 2 in Finocchiaro et al. (2019)), which corresponds precisely to the Max-Min-Margin loss S_{MM} defined in Eq. (8). In-

deed, S_{MM} and L have the same Bayes risk as

$$\begin{aligned} H_{\text{MM}}(q) &= \min_{v \in \mathbb{R}^k} \left(\max_{p \in \Delta} \min_{y \in \mathcal{Y}} L_y^\top p + v^\top p \right) - v^\top q \\ &= \min_{v \in \mathbb{R}^k} (-H_L)^*(v) - v^\top q = H_L(q), \end{aligned}$$

where $h^*(u) = \sup_{s \in \mathbb{R}^k} u^\top s - h(s)$ is the Fenchel conjugate of h . It can be checked (Nowak-Vila et al., 2020) that the embedding is $\psi(y) = -L_y$ and it is always Fisher consistent to L under the argmax decoding.

3.2 Preliminary Results

Relationship between losses and Bayes risks.

What makes the Max-Min-Margin loss simple to analyze is its Fenchel-Young structure (Blondel et al., 2020), i.e., it can be written in the form $S(v, y) = \Omega^*(v) - v_y$, for a certain convex function Ω defined in the simplex. We extend this notion by allowing the convex function Ω to depend on the label y as

$$S(v, y) = (\Omega^y)^*(v) - v_y. \quad (10)$$

The losses S_{M} , S_{RM} and S_{MM} can be written in this form with

$$\begin{aligned} \Omega_{\text{MM}}^y(q) &= - \min_{y' \in \mathcal{Y}} L_{y'}^\top q + i_\Delta(q), \\ \Omega_{\text{M}}^y(q) &= -L_y^\top q + i_\Delta(q), \\ \Omega_{\text{RM}}^y(q) &= -L_y^\top q + i_{\Delta(y)}(q), \end{aligned}$$

where $i_U(u) = 0$ if $u \in U$ and ∞ otherwise. The first equation is the only one independent of y and we remove its dependence by simply writing Ω_{MM} . The following Prop. 3.3 relates the functions Ω_{MM} , Ω_{M}^y and Ω_{RM}^y .

Proposition 3.3. *The following holds: $\Omega_{\text{MM}} = \max_{y \in \mathcal{Y}} \Omega_{\text{M}}^y$ and $(\Omega_{\text{MM}})^* = \max_{y \in \mathcal{Y}} (\Omega_{\text{RM}}^y)^*$.*

Proof. The first identity is trivial from the definition. For the second identity, note that by construction the prediction sets necessarily cover the simplex as $\Delta = \cup_{y' \in \mathcal{Y}} \Delta(y')$. Hence,

$$\begin{aligned} (\Omega_{\text{MM}})^*(v) &= \max_{y \in \mathcal{Y}} \left(\max_{q \in \Delta(y)} \min_{y' \in \mathcal{Y}} L_{y'}^\top q + v^\top q \right) \\ &= \max_{y \in \mathcal{Y}} \left(\max_{q \in \Delta(y)} L_y^\top q + v^\top q \right) \\ &= \max_{y \in \mathcal{Y}} (\Omega_{\text{RM}}^y)^*(v), \end{aligned}$$

where we have used the fact that $\min_{y' \in \mathcal{Y}} L_{y'}^\top q = L_y^\top q$ whenever $q \in \Delta(y)$ by construction. \square

Moreover, it can be readily seen that $S_{\text{RM}}(v, y) \leq S_{\text{MM}}(v, y) \leq S_{\text{M}}(v, y)$, for all $v \in \mathbb{R}^k$ and $y \in \mathcal{Y}$. See Fig. 1 (left and middle) and Fig. 2 for the shape of

these losses when L is the 0-1 loss for two and three dimensions, respectively. Our main quantity of interest is the Bayes risk H_S , because by comparing it with H_L we are able to tell whether L is embedded by S using Prop. 3.2, thus proving consistency. The following Prop. 3.4 gives the form of the Bayes risks.

Proposition 3.4 (Bayes risks). *For all $q \in \Delta$, the Bayes risks read*

$$\begin{aligned} H_{\text{MM}}(q) &= H_L(q), \\ H_{\text{M}}(q) &= \max_{Q \in U(q,q)} \langle L, Q \rangle_{\text{F}}, \\ H_{\text{RM}}(q) &= \max_{Q \in U(q,q) \cap \mathcal{C}_L} \langle L, Q \rangle_{\text{F}}, \end{aligned}$$

where $U(q, q) = \{Q \in \mathbb{R}^{k \times k} \mid Q\mathbf{1} = q, Q^\top \mathbf{1} = q, Q \succeq 0\}$, $\mathcal{C}_L = \{Q \in \mathbb{R}^{k \times k} \mid (1L_y^\top - L)Q_y \preceq 0, \forall y \in \mathcal{Y}\}$ and $\langle \cdot, \cdot \rangle_{\text{F}}$ denotes the Frobenius inner product. Moreover, we have that $H_{\text{RM}} \leq H_{\text{MM}} \leq H_{\text{M}}$, and there exists L for which $H_L \neq H_{\text{M}}$ and/or $H_{\text{RM}} \neq H_L$.

The proof can be found in Section 1 of Appendix. As a corollary, the only one of these losses always embedding L is the Max-Min-Margin loss. Our sufficient conditions for consistency will correspond to the conditions on L for which the Bayes risks H_{M} and/or H_{RM} are equal (or proportional) to H_L , thus implying consistency. In the next section, we use the expressions of the Bayes risks given by Prop. 3.4 as a basis to prove the consistency results.

4 FISHER CONSISTENCY ANALYSIS

4.1 Analysis of Max-Margin Loss

In this section we want to provide a necessary condition for consistency of the Max-Margin loss (Thm. 2.1). Note that it is not enough to provide a condition for which $H_S \neq H_L$, as S consistent to L does not imply S embedding L . The following Prop. 4.1 gives a necessary condition for consistency in terms of the extreme points of the prediction sets and the Bayes risk of S .

Proposition 4.1. *Assume S is Fisher consistent to L . Then, for any extreme point $q \in \Delta$ of a prediction set $\Delta(y)$, $y \in \mathcal{Y}$, we necessarily have $\{q\} = \partial(-H_S)^*(v)$ for some $v \in \mathbb{R}^k$.*

The set $\partial h(x)$ denotes the subgradient of the function h . This result is proven in Section 2 of Appendix. To use this necessary condition, we first have to compute the Fenchel conjugate $(-H_{\text{M}})^*$ of the Max-Margin loss. This is given by the following Prop. 4.2.

Proposition 4.2. *If L is symmetric, then*

$$(-H_{\text{M}})^*(v) = \max_{y, y' \in \mathcal{Y}} L(y, y') + \frac{v_y + v_{y'}}{2},$$

for all $v \in \mathbb{R}^k$.

As a corollary, we obtain the specific form of the images of the sub-gradient mapping $\partial(-H_{\text{M}})^*$ at the differentiable points, i.e., when the sub-gradient is a singleton.

Corollary 4.3. *The 0-dimensional images of $\partial(-H_{\text{M}})^*$ are of the form $q = \frac{1}{2}(e_y + e_{y'})$, $y, y' \in \mathcal{Y}$.*

Note that the points $\{q\} \in \text{Im } \partial(-H_{\text{M}})^*$ are independent of the discrete loss L . In Fig. 3 (left), we show that this is not the case for the extreme points of the prediction sets $\Delta(y)$'s of a generic L . In this example, the points $\frac{1}{2}(e_y + e_{y'})$'s are extreme points of the prediction sets because L is symmetric, but the extreme point in the interior is not of this form. By moving this point to $\frac{1}{2}(e_1 + e_3)$, we enforce $L(1, 2) + L(2, 3) = L(1, 3)$, which corresponds precisely to the necessary condition given by Thm. 2.1 when $z = y_2$. The generic necessary condition is obtained by extending this argument to consider all possibilities in larger output spaces. The full proof can be found in Section 2 of Appendix.

Consistency of distances defined in trees. This result is proved by showing $H_{\text{M}} = 2H_L$ ($H_S \propto H_L$ also implies consistency (Finocchiaro et al., 2019)) whenever L is a distance defined in a tree, and it is done by proving equality of the Fenchel conjugates $(-H_{\text{M}})^* = (-2H_L)^*$. The proof is based on the analysis of the extreme points of a certain polytope defined in terms of L , and can be found in Section 2 of Appendix.

Partial positive results on consistency. In this section, we generalize the well-known result by Liu (2007) which states that the Max-Margin loss is consistent to the 0-1 loss under the dominant label condition. More specifically, we provide a generalization of this result to all distance losses, i.e., symmetric and satisfying the triangle inequality.

Proposition 4.4. *If L is a distance, then we have $H_{\text{M}} \leq 2H_L$ and $H_{\text{M}}(q) = 2H_L(q)$ (S embeds $2L$, thus consistent), for all $q \in \Delta$ such that $\max_{y \in \mathcal{Y}} q_y \geq 1/2$.*

Some intuition can be obtained for $k = 3$. The interior of the set delimited by the dashed lines in the left image from Fig. 3 shows precisely the points where $\max_{y \in \mathcal{Y}} q_y \leq 1/2$. If L is a distance, then the interior extreme point can only move inside this region, and hence, it remains consistent at the exterior of this set, which is precisely where the dominant label condition is satisfied. The proof of this result can be found in Section 2 of Appendix.




Figure 3: **Left:** Prediction sets $\Delta(y)$ for a symmetric loss. The points given by Cor. 4.3 are the 3 extremes of the simplex and the middle points in the faces of the simplex. The interior point is not of the form $\frac{1}{2}(e_y + e_{y'})$ for any $y, y' \in \{1, 2, 3\}$. **Middle:** The sufficient condition from Prop. 4.6 is satisfied for any point in the simplex. **Right:** The sufficient condition from Prop. 4.6 is not satisfied as the prediction set $\Delta(2)$ (in green) intersects the line $q_2 = 0$. Moreover, inconsistency can be shown in this case by comparing the affine regions of the Bayes risk H_{RM} with these predictions sets.

4.2 Analysis of Restricted-Max-Margin Loss

In this section we sketch the proof of the consistency of the Restricted-Max-Margin loss (Thm. 2.4). We prove that S_{RM} embeds L by showing equal Bayes risks using the expressions from Prop. 3.4. The proof is done in two steps. In the first one (Prop. 4.5), we show that under a condition on the extreme points of the prediction sets both Bayes risks are equal, while in the second we obtain that this condition is satisfied whenever the hypothesis from Thm. 2.4 holds.

Proposition 4.5. *If $qq^\top \in \mathcal{C}_L = \{Q \in \mathbb{R}^{k \times k} \mid (1L_y^\top - L)Q_y \preceq 0, \forall y \in \mathcal{Y}\}$ for all extreme points $q \in \Delta$ of the prediction sets $\Delta(y)$'s, then $H_{\text{RM}} = H_L$.*

Proof. By Prop. 3.4, we already know that $H_{\text{RM}} \leq H_L$, so we only need to show that $H_{\text{RM}} \geq H_L$. If q is an extreme point of the prediction set $\Delta(y)$, then $H_{\text{RM}}(q) \geq H_L(q)$ by taking the matrix $qq^\top \in U(q, q) \cap \mathcal{C}_L$, which satisfies $\langle L, qq^\top \rangle_{\text{F}} = \sum_{y'} q_{y'} L_{y'}^\top q \geq \sum_{y'} q_{y'} L_y^\top q = L_y^\top q = H_L(q)$, where we have used that $L_{y'}^\top q \geq L_y^\top q$ for all $y' \in \mathcal{Y}$ as $q \in \Delta(y)$. If q is not an extreme point, then it can be written as a convex combination of extreme points q_i of $\Delta(y)$ as $q = \sum_{i=1}^m \alpha_i q_i$ with $\alpha^\top \mathbf{1} = 1, \alpha \succeq 0$. Then, it is straightforward to show that the matrix $Q = \sum_{i=1}^m \alpha_i q_i q_i^\top$ is in $U(q, q) \cap \mathcal{C}_L$, and satisfies $\langle L, Q \rangle_{\text{F}} \geq L_y^\top q = H_L(q)$. More details in Section 3 of Appendix. \square

Proposition 4.6. *Assume that for every $q \in \Delta$, if y is optimal for q (i.e., $q \in \Delta(y)$), then $q_y > 0$. In this case, all extreme points of the prediction sets satisfy $qq^\top \in \mathcal{C}_L$.*

The proof of this result and the fact that consistency of S_{M} implies consistency of S_{RM} can be found in Section 3 of Appendix. In Fig. 3 (middle and right), we show

geometrically in dimension $k = 3$ what this condition means.

4.3 Argmax Decoding

The positive consistency results of both S_{M} and S_{RM} do not specify whether the argmax decoding $d(v) = \arg \max_{y \in \mathcal{Y}} v_y$ is consistent, but just that there exists a decoding for which consistency holds. From Nowak-Vila et al. (2020), we know that the set of minimizers of the Max-Min-Margin loss S_{MM} is

$$v_{\text{MM}}^*(q) = -\text{hull}(L_y)_{y \in y^*(q)} + \mathcal{N}_\Delta(q), \quad (11)$$

for all $q \in \Delta$, where hull denotes the convex hull and $\mathcal{N}_\Delta(q) = \{u \in \mathbb{R}^k \mid u^\top(p - q) \leq 0, \forall p \in \Delta\}$ is the normal cone of the simplex at the point q . In this case, the argmax decoding is consistent because $\arg \max_y v_y \in y^*(q)$ whenever $v \in v_{\text{MM}}^*(q)$ if Eq. (11) is satisfied. The following Thm. 4.7 shows that the same holds true for the other losses whenever they embed L (or $2L$).

Theorem 4.7. *S_{M} and S_{RM} are consistent to L under the argmax decoding whenever $H_{\text{M}} = 2H_L$ (S_{M} embeds $2L$) and $H_{\text{RM}} = H_L$ (S_{RM} embeds L), respectively.*

5 CONCLUSION AND FUTURE DIRECTIONS

In this work, we have analyzed the consistency properties of the well-known Max-Margin loss for general classification tasks. We show that a restrictive condition on the task, which is not satisfied by most of the used losses in practice, is necessary for Fisher consistency. We also show consistency for the absolute deviation loss used in ordinal regression, and more generally to any distance defined in a tree, showing that the

necessary condition is meaningful. Moreover, we have overcome this limitation and introduced a novel generalization of the binary SVM loss called Restricted-Max-Margin loss, which maintains the maximization over the loss-augmented scores and it is consistent under milder conditions on the task at hand. Two important questions remain unanswered: first, whether the proposed necessary condition for the Max-Margin is sufficient or not, and secondly, whether there exist tractable linear maximization oracles over the prediction sets $\Delta(y)$ for specific structured prediction problems, which would make the Restricted-Max-Margin loss more attractive than the Max-Min-Margin loss in practice.

Acknowledgements

This work was funded in part by the French government under management of Agence Nationale de la Recherche as part of the ‘Investissements d’avenir’ program, reference ANR-19-P3IA-0001 (PRAIRIE 3IA Institute). We acknowledge support of the European Research Council (grant SEQUOIA 724063). We also acknowledge support of the European Research Council (grant REAL 947908).

References

- Bernhard E Boser, Isabelle M Guyon, and Vladimir N Vapnik. A training algorithm for optimal margin classifiers. In *Proceedings of the fifth annual workshop on Computational learning theory*, pages 144–152, 1992.
- Corinna Cortes and Vladimir Vapnik. Support-vector networks. *Machine Learning*, 20(3):273–297, 1995.
- Koby Crammer and Yoram Singer. On the algorithmic implementation of multiclass kernel-based vector machines. *Journal of Machine Learning Research*, 2(Dec):265–292, 2001.
- Ben Taskar, Carlos Guestrin, and Daphne Koller. Max-margin Markov networks. In *Advances in Neural Information Processing Systems*, pages 25–32, 2004.
- Ioannis Tsochantaridis, Thorsten Joachims, Thomas Hofmann, and Yasemin Altun. Large margin methods for structured and interdependent output variables. *Journal of Machine Learning Research*, 6 (Sep):1453–1484, 2005.
- Yufeng Liu. Fisher consistency of multicategory support vector machines. In *Artificial Intelligence and Statistics*, pages 291–298, 2007.
- Ben Taskar, Simon Lacoste-Julien, and Dan Klein. A discriminative matching approach to word alignment. In *Proceedings of Human Language Technology Conference and Conference on Empirical Methods in Natural Language Processing*, pages 73–80, 2005.
- Tibério S. Caetano, Julian J. McAuley, Li Cheng, Quoc V. Le, and Alex J. Smola. Learning graph matching. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 31(6):1048–1058, 2009.
- Noah A. Smith. Linguistic structure prediction. *Synthesis Lectures on Human Language Technologies*, 4 (2):1–274, 2011.
- Rizal Fathony, Anqi Liu, Kaiser Asif, and Brian Ziebart. Adversarial multiclass classification: A risk minimization perspective. In *Advances in Neural Information Processing Systems*, pages 559–567, 2016.
- John Duchi, Khashayar Khosravi, and Feng Ruan. Multiclass classification, information, divergence and surrogate risk. *The Annals of Statistics*, 46(6B): 3246–3275, 2018.
- Alex Nowak-Vila, Francis Bach, and Alessandro Rudi. Consistent structured prediction with max-min margin markov networks. In *Proceedings of the International Conference on Machine Learning (ICML)*, 2020.
- Nachman Aronszajn. Theory of reproducing kernels. *Transactions of the American Mathematical Society*, 68(3):337–404, 1950.
- Vladimir Vapnik. Principles of risk minimization for learning theory. In *Advances in Neural Information Processing Systems*, pages 831–838, 1992.
- Vladimir Vapnik. *The nature of statistical learning theory*. Springer Science & Business Media, 2013.
- Peter L. Bartlett, Michael I. Jordan, and Jon D. McAuliffe. Convexity, classification, and risk bounds. *Journal of the American Statistical Association*, 101(473):138–156, 2006.
- Luc Devroye, László Györfi, and Gábor Lugosi. *A Probabilistic Theory of Pattern Recognition*, volume 31. Springer Science & Business Media, 2013.
- Thorsten Joachims. Training linear SVMs in linear time. In *Proceedings of the SIGKDD International Conference on Knowledge Discovery and Data Mining*, pages 217–226. ACM, 2006.
- Gökhan Bakır, Thomas Hofmann, Bernhard Schölkopf, Alexander J. Smola, Ben Taskar, and S.V.N. Vishwanathan. *Predicting Structured Data*. MIT press, 2007.
- Harish G. Ramaswamy, Ambuj Tewari, and Shivani Agarwal. Consistent algorithms for multiclass classification with an abstain option. *Electronic Journal of Statistics*, 12(1):530–554, 2018.

- David McAllester. Generalization bounds and consistency. *Predicting Structured Data*, pages 247–261, 2007.
- Ürün Dogan, Tobias Glasmachers, and Christian Igel. A unified view on multi-class support vector classification. *J. Mach. Learn. Res.*, 17(45):1–32, 2016.
- Jason Weston and Chris Watkins. Support vector machines for multi-class pattern recognition. In *ESANN*, volume 99, pages 219–224, 1999.
- Yoonkyung Lee, Yi Lin, and Grace Wahba. Multi-category support vector machines: Theory and application to the classification of microarray data and satellite radiance data. *Journal of the American Statistical Association*, 99(465):67–81, 2004.
- Youssef Mroueh, Tomaso Poggio, Lorenzo Rosasco, and Jean-Jacques Slotine. Multiclass learning with simplex coding. In *Advances in Neural Information Processing Systems*, pages 2789–2797, 2012.
- James Petterson, Jin Yu, Julian J. McAuley, and Tibério S. Caetano. Exponential family graph matching and ranking. In *Advances in Neural Information Processing Systems*, pages 1455–1463, 2009.
- Ambuj Tewari and Peter L. Bartlett. On the consistency of multiclass classification methods. *Journal of Machine Learning Research*, 8(May):1007–1025, 2007.
- Jessica Finocchiaro, Rafael Frongillo, and Bo Waggoner. An embedding framework for consistent polyhedral surrogates. In *Advances in Neural Information Processing Systems*, 2019.
- Mathieu Blondel, André FT Martins, and Vlad Niculae. Learning with Fenchel-Young losses. *Journal of Machine Learning Research*, 21(35):1–69, 2020.
- R Tyrrell Rockafellar. *Convex analysis*. princeton landmarks in mathematics, 1997.
- Niclas Andreasson, Michael Patriksson, and Anton Evgrafov. *An Introduction to Continuous Optimization: Foundations and Fundamental Algorithms*. Courier Dover Publications, 2020.

Supplementary Material: On the Consistency of Max-Margin Losses

Outline. The supplementary material is organized as follows. In Appendix A, we prove general results on embeddings of losses, we compute the Bayes risks for each of the losses and we provide an algebraic characterization of the extreme points of the prediction sets. In Appendix B and Appendix C, we provide the main results of the Max-Margin loss and the Restricted-Max-Margin loss, respectively.

A PRELIMINARY RESULTS

A.1 Results on Embeddability of Losses

Proposition A.1. *Let $\psi : \mathcal{Y} \rightarrow \mathbb{R}^k$ be an embedding of the output space. If $H_S = H_L$ and $S(\psi(y)) = L_y$ for all $y \in \mathcal{Y}$, then S embeds L with embedding ψ .*

Proof. To prove that S embeds L with embedding $\psi(y) = -L_y$, we need to show that

$$y \in y^*(q) \iff \varphi(y) \in v^*(q).$$

If $y \in y^*(q)$, then

$$H_L(q) = L_y^\top q = S(\varphi(y))^\top q = H(q) = \min_{v \in \mathbb{R}^k} S(v)^\top q.$$

Thus, $S(\varphi(y))^\top q = \min_{v \in \mathbb{R}^k} S(v)^\top q$ implies that necessarily $\varphi(y) \in v^*(q)$. Similarly, if $\varphi(y) \in v^*(q)$, then $\min_z L_z^\top q = L_y^\top q$ which implies $y \in y^*(q)$. \square

A.2 Bayes risk identities

The following Lemma A.2 provides an identity which will be useful to provide the forms of the Bayes risk for S_M and S_{RM} .

Lemma A.2. *Let $\mathcal{C}_y \subseteq \Delta$, $\Omega^y(q) = -L_y^\top q + i_{\mathcal{C}_y}(q)$ and $S(v, y) = (\Omega^y)^*(v) - v_y$ for every $y \in \mathcal{Y}$. Then,*

$$H_S(q) = \max_{\substack{\sum_y q_y \nu_y = q \\ \nu_y \in \mathcal{C}_y}} \sum_y q_y L_y^\top \nu_y. \quad (12)$$

Proof. Recall the definition of the Bayes risk $H(q) = \min_{v \in \mathbb{R}^k} S(v)^\top q$. Using the structural assumption on S , we can re-write it as

$$H(q) = \min_{v \in \mathbb{R}^k} \sum_{y \in \mathcal{Y}} q_y (\Omega^y)^*(v) - v^\top q = - \max_{v \in \mathbb{R}^k} v^\top q - \sum_{y \in \mathcal{Y}} q_y (\Omega^y)^*(v) = - \left(\sum_{y \in \mathcal{Y}} q_y (\Omega^y)^* \right)^*(q).$$

Recall that if the functions h_i are convex, then the conjugate of the sum is the infimum convolution of the individual conjugates (Rockafellar, 1997) as

$$\left(\sum_i h_i \right)^*(t) = \min_{\sum_i x_i = t} \sum_i h_i^*(x_i).$$

If we apply this property to the functions $h_i = q_i(\Omega^y)^*$, we obtain:

$$\begin{aligned}
 -\left(\sum_{y \in \mathcal{Y}} q_y(\Omega^y)^*\right)^*(q) &= -\min_{\sum_{y \in \mathcal{Y}} \nu_y = q} \sum_{y \in \mathcal{Y}} (q_y(\Omega^y)^*)^*(\nu_y) \\
 &= -\min_{\sum_{y \in \mathcal{Y}} \nu_y = q} \sum_{y \in \mathcal{Y}} q_y \Omega^y(\nu_y/q_y) && (ah)^*(x) = ah^*(x/a) \\
 &= -\min_{\substack{\sum_{y \in \mathcal{Y}} \nu_y = q \\ \nu_y/q_y \in \mathcal{C}_y, \forall y \in \mathcal{Y}}} -\sum_{y \in \mathcal{Y}} L_y^\top \nu_y && \Omega^y(q) = -L_y^\top q + i_{\mathcal{C}_y}(q) \\
 &= \max_{\substack{\sum_{y \in \mathcal{Y}} \nu_y = q \\ \nu_y/q_y \in \mathcal{C}_y, \forall y \in \mathcal{Y}}} \sum_{y \in \mathcal{Y}} L_y^\top \nu_y && \text{redefine } \nu_y \text{ as } \nu_y/q_y \\
 &= \max_{\substack{\sum_{y \in \mathcal{Y}} q_y \nu_y = q \\ \nu_y \in \mathcal{C}_y, \forall y \in \mathcal{Y}}} \sum_{y \in \mathcal{Y}} q_y L_y^\top \nu_y.
 \end{aligned}$$

□

The following Prop. A.3 provides us with the first part of Proposition 3.4.

Proposition A.3 (Bayes risks). *For all $q \in \Delta$, the Bayes risks read*

$$\begin{aligned}
 H_{\text{MM}}(q) &= \min_{y \in \mathcal{Y}} L_y^\top q = H_L(q) \\
 H_{\text{M}}(q) &= \max_{Q \in U(q, q)} \langle L, Q \rangle_{\text{F}} \\
 H_{\text{RM}}(q) &= \max_{Q \in U(q, q) \cap \mathcal{C}_L} \langle L, Q \rangle_{\text{F}},
 \end{aligned}$$

where

$$U(q, q) = \{Q \in \mathbb{R}_+^{k \times k} \mid Q\mathbf{1} = q, Q^\top \mathbf{1} = q\}, \quad \text{and} \quad \mathcal{C}_L = \{Q \in \mathbb{R}^{k \times k} \mid (1L_y^\top - L)Q_y \leq 0, \forall y \in \mathcal{Y}\}.$$

Proof. The first identity is trivial and has already been derived in the main body of the paper. We use the above Lemma A.2 to obtain the identities corresponding to S_{M} and S_{RM} .

1. *Bayes risk of Max-Margin:* In this case $\mathcal{C}_y = \Delta$. If we define $\Gamma \in \mathbb{R}^{k \times k}$ as the matrix whose rows are ν_y , the maximization reads

$$\max_{\substack{\Gamma^\top q = q \\ \Gamma \mathbf{1} = \mathbf{1} \\ \Gamma \geq 0}} \sum_{y \in \mathcal{Y}} q_y L_y^\top \Gamma_y$$

If we now define $Q \in \mathbb{R}^{k \times k}$ as $Q = \text{diag}(q)\Gamma$, i.e., $Q_y = q_y \Gamma_y$, the objective can be re-written as a matrix scalar product as

$$\sum_{y \in \mathcal{Y}} q_y L_y^\top \Gamma_y = \sum_{y \in \mathcal{Y}} L_y^\top (q_y \Gamma_y) = \sum_{y \in \mathcal{Y}} L_y^\top Q_y = \langle L, Q \rangle_{\text{F}}.$$

Whenever $q_y > 0$, the change of variables $Q_y = q_y \Gamma_y$ is invertible and the constraints satisfy

$$(Q^\top \mathbf{1})_y = q_y \iff (\Gamma^\top q)_y = q_y \tag{13}$$

$$(Q\mathbf{1})_y = q_y \iff (\Gamma\mathbf{1})_y = 1 \tag{14}$$

$$Q_y \geq 0 \iff \Gamma_y \geq 0. \tag{15}$$

On the other hand, if $q_y = 0$ then $Q_y = 0$ but the objective is not affected as it is independent of Γ_y .

2. *Bayes risk of Restricted-Max-Margin:* In this case $\mathcal{C}_y = \Delta(y) = \{q \in \Delta \mid (L_y - L_{y'})^\top q \leq 0, \forall y' \in \mathcal{Y}\}$. The maximization now reads

$$\max_{\substack{\Gamma^\top q = q \\ \Gamma \mathbf{1} = \mathbf{1} \\ \Gamma \geq 0 \\ (1L_y^\top - L)\Gamma_y \leq 0, \forall y}} \sum_{y \in \mathcal{Y}} q_y L_y^\top \Gamma_y.$$

The result follows as $(1L_y^\top - L)\Gamma_y \leq 0$ if and only if $(1L_y^\top - L)Q_y \leq 0$ whenever $q_y > 0$.

□

A.3 Extreme points of a polytope

We will need to analyse the extreme points of the polytope $\mathcal{P} = \{(q, u) \in \mathbb{R}^{k+1} \mid q \in \Delta, L_y^\top q \geq u, \forall y \in \mathcal{Y}\} \subseteq \mathbb{R}^{k+1}$ in the proof of the sufficient condition for consistency of Max-Margin in Appendix B.3.

Algebraic characterization of extreme points of a polyhedron. The following Prop. A.4 provides us with an algebraic characterization of the extreme points of a polyhedron $\mathcal{Q} = \{x \in \mathbb{R}^n \mid Ax \succeq b\}$.

Proposition A.4 (Theorem 3.17 in Andreasson et al. (2020)). *Let $x \in \mathcal{Q} = \{x \in \mathbb{R}^n \mid Ax \succeq b\}$, where $A \in \mathbb{R}^{m \times n}$ has $\text{rank}(A) = n$ and $b \in \mathbb{R}^m$. Let $I \subseteq [m]$ be a set of indexes for which the subsystem is an equality, i.e., $A_I x_I = b_I$ with $Ax_I \succeq b$. Then x_I is an extreme point of \mathcal{Q} if and only if $\text{rank}(A_S) = n$.*

Let $\mathcal{P} \subseteq \mathbb{R}^{k+1}$ be the polyhedron defined as

$$\mathcal{P} = \{(q, u) \in \mathbb{R}^{k+1} \mid q \in \Delta, L_y^\top q \geq u, \forall y \in \mathcal{Y}\} \subseteq \mathbb{R}^{k+1}.$$

The polyhedron \mathcal{P} can be written as $\mathcal{P} = \{x = (q, u) \in \mathbb{R}^{k+1} \mid Ax \succeq b\}$ where

$$\begin{array}{c} S \\ T \end{array} \left\{ \begin{array}{c} \left(\begin{array}{c|c} \mathbf{L} & \begin{array}{c} \vdots \\ -1 \\ \vdots \end{array} \\ \hline \mathbf{Id} & \begin{array}{c} \vdots \\ 0 \\ \vdots \end{array} \\ \hline \dots & \begin{array}{c} 1 \\ \dots \\ -1 \end{array} \\ \dots & \begin{array}{c} \dots \\ -1 \\ \dots \end{array} \\ \hline & \begin{array}{c} 0 \\ 0 \end{array} \end{array} \right) \begin{pmatrix} q \\ u \end{pmatrix} \succeq \begin{pmatrix} \vdots \\ 0 \\ \vdots \\ 0 \\ \vdots \\ 1 \\ -1 \end{pmatrix} \end{array} \right. \quad (16)$$

with $A \in \mathbb{R}^{(2k+2) \times (k+1)}$ and $b \in \mathbb{R}^{k+1}$. Note that $\text{rank}(A) = k + 1$. Given $x = (q, u)$, define $S, T \subseteq \mathcal{Y}$ as the subsets of outputs such that

$$y \in S \iff L_y^\top q = u, \quad y \in T \iff q_y = 0, \quad (17)$$

i.e., S and T correspond to the indexes of the first and second block of the matrix A for which the inequality holds as an equality, respectively. More concretely, if I are the indices of A for which $A_I x_I = b_I$, we have that $I = S \cup (k+T) \cup \{2k+1\} \cup \{2k+1\}$, because the last two inequalities must be an equality as $q \in \Delta$. Moreover, the sets S have the following properties:

- We necessarily have $|S| \geq 1$: if $S = \emptyset$, then $\text{rank}(A_I) = k$ and so the rank is not maximal, thus $x = (q, u)$ cannot be an extreme point.
- We necessarily have $|S| + |T| \geq k$ (using the fact that $\text{rank}(A) = k + 1$).

B RESULTS ON MAX-MARGIN LOSS

B.1 Bayes Risk of Max-Margin for Symmetric Losses

The following Prop. B.1 gives another expression of the Bayes risk of S_M and its Fenchel conjugate assuming the loss L is symmetric.

Proposition B.1. *Let $H_M(q) = \max_{Q \in U(q,q)} \langle L, Q \rangle_F$ and assume L symmetric. Then, the following identities hold:*

$$H_M(q) = \min_{\substack{Q \succeq 0 \\ \frac{1}{2}(a_y + a_{y'}) \geq L(y,y')}} a^\top q, \quad \forall q \in \Delta, \quad (18)$$

$$(-H_M)^*(v) = \max_{y,y' \in \mathcal{Y}} L(y,y') + \frac{v_y + v_{y'}}{2}, \quad \forall v \in \mathbb{R}^k. \quad (19)$$

Proof. The first part corresponds to the dual of the maximization problem defining the Bayes risk H_M when L is symmetric:

$$\begin{aligned} - \min_{Q \in U(q,q)} -\langle L, Q \rangle_F &= \min_{Q \succeq 0} \max_{a,b \in \mathbb{R}^k} a^\top (Q1 - q) + b^\top (Q^\top 1 - q) - \langle L, Q \rangle_F \\ &= - \max_{a,b \in \mathbb{R}^k} -a^\top q - b^\top q + \min_{Q \succeq 0} a^\top Q1 + b^\top Q^\top 1 - \langle L, Q \rangle_F \end{aligned}$$

We can now re-write the minimization objective as a matrix scalar product with Q as $a^\top Q1 = \text{Tr}(a^\top Q1) = \text{Tr}(Q1a^\top) = \langle Q, a1^\top \rangle_F$ and analogously $b^\top Q1 = \langle Q, 1b^\top \rangle_F$. Hence, the objective of the minimum becomes $\langle a1^\top + 1b^\top - L, Q \rangle_F$, which gives

$$\min_{Q \succeq 0} \langle a1^\top + 1b^\top - L, Q \rangle_F = \begin{cases} 0 & \text{if } a1^\top + 1b^\top - L \succeq 0 \\ -\infty & \text{otherwise} \end{cases}.$$

We obtain the following minimization problem in $a, b \in \mathbb{R}^k$

$$- \max_{a1^\top + 1b^\top \succeq L} -(a+b)^\top q = \min_{a1^\top + 1b^\top \succeq L} (a+b)^\top q.$$

Using that L is symmetric, we can add the constraint $a = b$. In order to see this, let (a^*, b^*) be a solution of the linear problem. If L is symmetric, then (b^*, a^*) is also a solution, which implies that $\frac{1}{2}(a^* + b^*, a^* + b^*)$ too. Hence, we can assume $a = b$ and we obtain the desired result.

For the second part, note that if L is symmetric, the matrix Q can be assumed also symmetric. To see this, let $Q^* = \arg \max_{Q \in U(q,q)} \langle L, Q \rangle_F$. Then if L symmetric $(Q^*)^\top$ is also a solution, which means that $\frac{1}{2}(Q^* + (Q^*)^\top)$ too, which is symmetric. Hence, we can write

$$\begin{aligned} H_M(q) &= \max_{\substack{Q=Q^\top \\ Q1=q \\ Q \succeq 0}} \langle L, Q \rangle_F \\ &= \min_{v \in \mathbb{R}^k} \max_{\substack{Q=Q^\top \\ Q \in \text{Prob}(\mathcal{Y} \times \mathcal{Y})}} \langle L, Q \rangle_F - v^\top (Q1 - q) \\ &= \min_{v \in \mathbb{R}^k} \left\{ \underbrace{\max_{\substack{Q=Q^\top \\ Q \in \text{Prob}(\mathcal{Y} \times \mathcal{Y})}} \langle L + v1^\top, Q \rangle_F}_{(-H_M)^*(v)} \right\} - v^\top q, \end{aligned}$$

where at the last step we have used that $q \in \Delta$ and so $1^\top Q1 = 1$, which together with $Q \succeq 0$ implies $Q \in \text{Prob}(\mathcal{Y} \times \mathcal{Y})$. The extreme points of the problem domain $\{Q \in \text{Prob}(\mathcal{Y} \times \mathcal{Y})\}$ where the maximization of the linear objective is achieved are precisely the points $\{\frac{1}{2}(e_y + e_{y'})\}_{y,y' \in \mathcal{Y}}$. \square

B.2 Necessary Conditions for Consistency

Recall that S is consistent to L if there exists a decoding $d : \mathbb{R}^k \rightarrow \mathcal{Y}$ such that if $v \in v^*(q)$, then necessarily $d(v) \in y^*(q)$ for all $q \in \Delta$. A necessary condition for this to hold is that every level set of v^* must be included in a level set of y^* , which are precisely the prediction sets.

Lemma B.2. *If S is consistent to L , then for every $v \in \mathbb{R}^k$ there must exist a $y \in \mathcal{Y}$ such that*

$$(v^*)^{-1}(v) \subseteq (y^*)^{-1}(y) = \Delta(y). \quad (20)$$

Proof. If (20) does not hold, then there exists $q_1, q_2 \in (v^*)^{-1}(v)$ with $y^*(q_1) \cap y^*(q_2) = \emptyset$. However, Fisher consistency means that $v \in v^*(q_1)$ implies $d(v) \in y^*(q_1)$ and $v \in v^*(q_2)$ implies $d(v) \in y^*(q_2)$, which is not possible because $y^*(q_1) \cap y^*(q_2) = \emptyset$. \square

The following Cor. B.4 re-writes (20) in terms of the Bayes risk H_S .

Proposition B.3. *The level sets of v^* are the image of $-\partial(-H_S)^* : \mathbb{R}^k \rightarrow 2^\Delta$, i.e.,*

$$\text{Im}((v^*)^{-1}) = \text{Im}(-\partial(-H_S)^*). \quad (21)$$

Proof. First of all, note that $-\partial(-H_S)^* = (\partial H_M)^{-1}$ (Rockafellar, 1997). We have

$$H_S(q) = \min_{v \in \mathbb{R}^k} S(v)^\top q + i_\Delta(q), \quad \partial H_S(q) = S(\bar{v}) + \langle 1 \rangle, \quad \bar{v} \in v^*(q).$$

Let's now prove the two inclusions.

(\subseteq): Let $Q \in \text{Im}((v^*)^{-1})$. This means that there exists $V \in \mathbb{R}^k$ such that $V = \arg \min_{v \in \mathbb{R}^k} S(v)^\top q$ for all $q \in Q$. If we define $T = S(V) + \langle 1 \rangle$, then $T = \partial H_M(q)$ for all $q \in Q$.

(\supseteq): Let $Q \in \text{Im}((\partial H_M)^{-1})$. This means that there exists T such that $T = \partial H_M(q)$ for all $q \in Q$. For every $q \in Q$, the set T can be written as $T = S(v^*(q)) + \langle 1 \rangle$. To show that $v^*(q) = v^*(q')$ for all $q, q' \in Q$, we need to show that if $S(v) = S(v') + c1$, $v \in v^*(q)$, $v' \in v^*(q')$ for some $q, q' \in Q$, then necessarily $c = 0$. This is because $S(v(q))^\top q' \geq S(v(q'))^\top q' \implies c \geq 0$ and $S(v(q))^\top q \leq S(v(q'))^\top q \implies c \leq 0$. \square

Corollary B.4 (Necessary condition for consistency). *If S is Fisher consistent to L , then for every $v \in \mathbb{R}^k$, there exists $y \in \mathcal{Y}$ such that*

$$-\partial(-H_S)^*(v) \subseteq \Delta(y). \quad (22)$$

Proof. This follows directly from Eq. (20) and Prop. B.3. \square

Proposition B.5 (Weaker necessary condition for consistency). *If S is consistent to L , then every extreme point of $\Delta(y)$ for some $y \in \mathcal{Y}$ must be a 0-dimensional image of $-\partial(-H_S)^*$.*

Proof. Let $\Delta_S(v) = -\partial(-H_S)^*(v)$. There exists a finite set $\mathcal{V} \subseteq \mathbb{R}^k$ such that $\bigcup_{v \in \mathcal{V}} \Delta_S(v) = \Delta(y)$. In particular, if q is an extreme point of $\Delta(y)$, then there exists $v \in \mathcal{V}$ such that $q \in \Delta_S(v)$. We need to show that q is also an extreme point of $\Delta_S(v)$. Indeed, if $\Delta_S(v) \subseteq \Delta(y)$ are polyhedrons and $q \in \Delta_S(v)$, $\Delta(y)$ is an extreme point of $\Delta(y)$, then it is also necessarily an extreme point of $\Delta_S(v)$. \square

Theorem B.6. *Let L be a symmetric loss with $k > 2$. If the Max-Margin loss is consistent to L , then L is a distance, and for every three outputs $y_1, y_2, y_3 \in \mathcal{Y}$, there exists $z \in \mathcal{Y}$ for which these the following three identities are satisfied:*

$$\begin{aligned} L(y_1, y_2) &= L(y_1, z) + L(z, y_2), \\ L(y_1, y_3) &= L(y_1, z) + L(z, y_3), \\ L(y_2, y_3) &= L(y_2, z) + L(z, y_3). \end{aligned}$$




Figure 4: These are the only possible possibilities of the prediction sets in a three-dimensional face of the simplex. The equations associated to each configuration is written below the corresponding simplex and all together can be compactly written as the necessary condition given by the theorem. An edge from a corner of the simplex to the middle point of the opposite side is not possible as $L(y, y') = 0$ if and only if $y = y'$ by assumption.

Proof. From Prop. B.1 and Prop. B.5, we obtain that if the Max-Margin loss is consistent to L , then the extreme points of the prediction sets $\Delta(y)$'s have to be of the form $1/2(e_y + e_{y'})$. Hence, the projection of the sets $\Delta(y)$'s into a three-dimensional simplex can only be of the form depicted in Fig. 4. The necessary condition follows directly from these possibilities (see caption of Fig. 4). Moreover, note that if the three identities of the theorem hold, then L is a distance. To see that, note that the triangle inequality holds for any triplet $y_1, y_2, y_3 \in \mathcal{Y}$ as:

$$\begin{aligned}
 L(y_1, y_2) &= L(y_1, z) + L(z, y_2) \\
 &= L(y_1, y_3) - L(z, y_3) + L(y_3, y_2) - L(y_3, z) \\
 &= L(y_1, y_3) + L(y_3, y_1) - 2L(y_3, z) \\
 &\leq L(y_1, y_3) + L(y_3, y_1).
 \end{aligned}$$

□

Examples of losses not satisfying the necessary condition. We now show that the examples exposed in Section 2 do not satisfy the necessary condition of Thm. B.6.

Lemma B.7. *Let L such with full rank loss matrix L and existing $q \in \text{int}(\Delta)$ for which all outputs optimal, i.e., $y^*(q) = \mathcal{Y}$. Then, the Max-Margin loss is not consistent to L .*

Proof. The point q , which is not of the form $1/2(e_y + e_{y'})$ for some $y, y' \in \mathcal{Y}$, is an extreme point of the polytope $\Delta(y) = \{q \in \Delta \mid L_z^\top q \geq L_y^\top q, \forall z \in \mathcal{Y}\}$ for every $y \in \mathcal{Y}$. This is because $q \in \Delta$ is the unique solution of $Lq = u$ with $u = L_y^\top q$ for all $y \in \mathcal{Y}$. Hence, by Prop. B.5, the Max-Margin loss is not consistent to L . □

Lemma B.8. *The Max-Margin loss is not consistent to the the Hamming loss on permutations $L(\sigma, \sigma') = \frac{1}{M} \sum_{m=1}^M 1(\sigma(m) \neq \sigma'(m))$ where σ, σ' permutations of size M .*

Proof. Take the transpositions $\sigma_1 = (3, 2), \sigma_2 = (2, 1), \sigma_3 = (3, 1)$. We have that $L(\sigma_i, \sigma_j) = 2/M$ for $i \neq j$ and $L(\sigma, \sigma') > \frac{2}{M}$ for all permutations $\sigma \neq \sigma'$. Hence, the necessary condition can't be satisfied. □

The Hamming loss with $M = k = 2$ is consistent and it is not defined in a tree. The Hamming loss $L(y, y') = \frac{1}{2}(1(y_1 \neq y'_1) + 1(y_2 \neq y'_2))$ is consistent as it decomposes additively and each term is consistent as it is the binary 0-1 loss. However, it can't be described as the shortest path distance in a tree, but rather the shortest path distance in a cycle of size four with all weights equal to $1/2$.




Figure 5: **Left:** The 4-partition of the output space where $y_1, y_2, y_3 \in S$. The different sets only communicate through the edges $z_i - y_i$ for $i = 1, 2, 3$ respectively. **Right:** The 3-partition of the output space where $s_1, s_2 \in S$ and the depicted path is the path between these two points.

B.3 Sufficient condition on the discrete loss L

Theorem B.9. *If L is a distance defined in a tree, then the Max-Margin loss S_M embeds L with embedding $\psi(y) = -L_y$ and it is consistent under the argmax decoding.*

Proof. We just have to prove that the extremes points of the polytope \mathcal{P} defined as

$$\mathcal{P} = \{(q, u) \in \mathbb{R}^{k+1} \mid q \in \Delta, L_y^\top q \geq u, \forall y \in \mathcal{Y}\} \subseteq \mathbb{R}^{k+1}$$

are of the form $(\frac{1}{2}(e_y + e_{y'}), \frac{1}{2}L(y, y'))$, where $y, y' \in \mathcal{Y}$. Using this, we obtain

$$\begin{aligned} (-H_{2L})^*(v) &= \max_{q \in \Delta} \min_{z \in \mathcal{Y}} 2L_z^\top q + v^\top q \\ &= \max_{(q, u) \in \mathcal{P}} 2u + v^\top q \\ &= \max_{(q, u) \in \text{ext}(\mathcal{P})} 2u + v^\top q \\ &= \max_{y, y' \in \mathcal{Y}} L(y, y') + \frac{v_y + v_{y'}}{2} = (-H_M)^*(v), \end{aligned}$$

for all $v \in \mathbb{R}^k$. This implies $H_{2L} = 2H_L = H_M$. We will use the algebraic framework introduced in Appendix A.3.

Let $x \in \text{ext}(\mathcal{P})$ and S are T the sets of indices for which $L_y^\top q = u$ and $q_{y'} = 0$ for $y \in S$ and $y' \in T$ (as defined in (17)).

If $|S| = 1$, then we necessarily have $|T| = k - 1$ because $|T| \geq k - |S| = k - 1$ and $|T| = k$ is not possible because q is in the simplex. In this case, the extreme point is equal to $q = (e_y, 0)$, which is of the desired form.

First part of the proof. If $|S| \geq 2$, let's prove that the elements in S must be necessarily aligned, i.e., contained in a chain (always true for $|S| = 2$). If we denote by $\text{SP}(s, s') \subseteq \mathcal{Y}$ the elements in the shortest path between s, s' , this means that there exists $s_1, s_2 \in S$ such that $s \in \text{SP}(s_1, s_2)$ for all $s \in S$.

If the elements in S are not aligned, then there exist pairwise different elements $y_1, y_2, y_3 \in S$ and $z_1, z_2, z_3 \in \mathcal{Y}$ (possibly repeated) such that the tree defining the loss L can be partitioned into four sets I, II, III, IV of the form depicted in the left Fig. 5, where the edges $y_i - z_i$ belong to the tree for $i = 1, 2, 3$.

$$\begin{aligned}
 L_{z_1}^\top q &= \sum_{y \in \mathcal{I}} L(z_1, y) q_y + \sum_{\substack{y \notin \mathcal{I} \\ y \neq z_1}} L(z_1, y) q_y \\
 &= \sum_{y \in \mathcal{I}} (L(z_1, y_1) + L(y_1, y)) q_y + \sum_{\substack{y \notin \mathcal{I} \\ y \neq z_1}} (L(y_1, y) - L(z_1, y_1)) q_y \\
 &= \sum_{y \neq z_1} L(y_1, y) q_y + \left(\sum_{y \in \mathcal{I}} q_i - \sum_{\substack{y \notin \mathcal{I} \\ i \neq z_1}} q_y \right) L(z_1, y_1) \\
 &= \sum_{y \in \mathcal{Y}} L(y_1, y) q_y + \left(\sum_{y \in \mathcal{I}} q_i - \sum_{y \notin \mathcal{I}} q_y \right) L(z_1, y_1) \\
 &= u + \left(\sum_{y \in \mathcal{I}} q_y - \sum_{y \notin \mathcal{I}} q_y \right) L(z_1, y_1).
 \end{aligned}$$

If we repeat the same procedure for II and III, we obtain that

$$\begin{cases}
 L_{z_1}^\top q &= u + \left(\sum_{y \in \mathcal{I}} q_y - \sum_{y \notin \mathcal{I}} q_y \right) L(z_1, y_1) \\
 L_{z_2}^\top q &= u + \left(\sum_{y \in \mathcal{II}} q_y - \sum_{y \notin \mathcal{II}} q_y \right) L(z_2, y_2) \\
 L_{z_3}^\top q &= u + \left(\sum_{y \in \mathcal{III}} q_y - \sum_{y \notin \mathcal{III}} q_y \right) L(z_3, y_3)
 \end{cases}$$

However, note that $\sum_{y \in \mathcal{I}} q_y + \sum_{y \in \mathcal{II}} q_y + \sum_{y \in \mathcal{III}} q_y + \sum_{y \in \mathcal{IV}} q_y = 1$, which implies that

$$\min_{\mathcal{A} \in \{\mathcal{I}, \mathcal{II}, \mathcal{III}\}} \left(\sum_{y \in \mathcal{A}} q_y - \sum_{y \notin \mathcal{A}} q_y \right) < 0.$$

Hence, there exists $i \in \{1, 2, 3\}$ for which $L_{z_i}^\top q < u$, which leads to a contradiction as $L_y^\top q \geq u$ for all $y \in \mathcal{Y}$.

Second part of the proof. Now that we have that the elements in S must be aligned, let's proceed with the proof by analyzing separately particular cases:

- ($S \cap T = \emptyset$): This means that $q_s > 0$ for all $s \in S$. Let $x = (\frac{1}{2}(e_{s_1} + e_{s_2}), \frac{1}{2}L(s_1, s_2))$, where $S \subseteq \text{SP}(s_1, s_2)$. Then, it satisfies the equality constraints as $L_s q = 1/2(L(s_1, s) + L(s, s_2)) = 1/2L(s_1, s_2)$ because $s \in \text{SP}(s_1, s_2)$ for all $s \in S$. Hence, it has to be equal to the unique solution of the linear system of equations.
- ($S \cap T \neq \emptyset$): Let's separate into two more cases.
 - ($\exists r_1, r_2 \in [k] \setminus T$ such that $S \subseteq \text{SP}(r_1, r_2)$): Let $x = (\frac{1}{2}(e_{r_1} + e_{r_2}), \frac{1}{2}L(r_1, r_2))$. Then, it satisfies the equality constraints as $L_s q = 1/2(L(r_1, s) + L(s, r_2)) = 1/2L(r_1, r_2)$ because $s \in \text{SP}(r_1, r_2)$ for all $s \in S$. Hence, it has to be equal to the unique solution of the linear system of equations.
 - ($\nexists r_1, r_2 \in [k] \setminus T$ such that $S \subseteq \text{SP}(r_1, r_2)$): We will show that this case is not possible. Consider the shortest path between s_1 and s_2 in S and the partition of the vertices of the tree into the sets I, II, III depicted in the right Fig. 5. We know that

$$\{\text{I} \cap ([k] \setminus T) = \emptyset\} \vee \{\text{II} \cap ([k] \setminus T) = \emptyset\}.$$

If this is not true, then taking $r_1 \in \text{I}$ and $r_2 \in \text{II}$ we obtain $S \subseteq \text{SP}(r_1, r_2)$. Assume that $\text{I} \cap ([k] \setminus T) = \emptyset$. We have that $L(s_1, y) > L(\bar{s}_1, y)$ for all $y \in [k] \setminus T$, which means that

$$L_{s_1}^\top q > L_{\bar{s}_1}^\top q.$$

This is a contradiction because $L_y^\top q \geq u = L_{s_1}^\top q$ as $s_1 \in S$. The case $\text{II} \cap ([k] \setminus T) = \emptyset$ can be done analogously.

Third part of the proof. By Prop. A.1, to prove that S_M embeds $2L$ with embedding $\psi(y) = -L_y$, we only need to show that $S_M(-L_y) = 2L_y$. For every $z \in \mathcal{Y}$, we have

$$\begin{aligned} S_M(-L_y, z) &= \max_{y' \in \mathcal{Y}} L(z, y') + (-L_y)^\top e_{y'} - (-L_y)^\top e_z \\ &= \max_{y' \in \mathcal{Y}} \{L(z, y') - L(y, y')\} + L(y, z) \\ &= 2L(y, z), \end{aligned}$$

where at the last step we have used $L(z, y') - L(y, y') \leq L(y, y')$ as L is a distance and so the maximization is achieved at $y' = y$.

Finally, the argmax decoding is consistent as it is an inverse of the embedding $\psi(y) = -L_y$ as

$$d(\psi(y)) = \arg \max_{y' \in \mathcal{Y}} -L(y, y') = \arg \min_{y' \in \mathcal{Y}} L(y, y') = y.$$

□

B.4 Partial Consistency through dominant label condition

Lemma B.10. *Let $q \in \Delta$ such that $q_1 \geq 1/2 \geq p_y$ for all $y \neq 1$. If L is a distance, then $L_1^\top q \leq L_y^\top q$ for all $y \in \mathcal{Y}$.*

Proof.

$$\begin{aligned} L_z^\top q &= q_1 L_{z1} + \sum_{y \neq 1, z} L_{zy} q_y \\ &\geq \frac{1}{2} L_{z1} + \sum_{y \neq 1, z} L_{zy} q_y \\ &= \left(\frac{1}{2} - \sum_{y \neq 1, z} q_y \right) L_{z1} + \sum_{y \neq 1, z} (L_{z1} + L_{zy}) q_y \\ &\geq \left(\frac{1}{2} - \sum_{y \neq 1, z} q_y \right) L_{z1} + \sum_{y \neq 1, z} L_{1y} q_y \\ &\geq L_{z1} q_z + \sum_{y \neq 1, z} L_{1y} q_y \\ &= \sum_{y \neq 1} L_{1y} q_y = L_1^\top q. \end{aligned}$$

□

Lemma B.11. *If L is a distance, then $H_M \leq 2H_L$.*

Proof. In particular, L is also symmetric. Recall that

$$H_L(q) = \min_{y \in \mathcal{Y}} L_y^\top q = \min_{\substack{a \succeq Lp, \\ p \in \Delta}} a^\top q = \min_{a \in \mathcal{P}_L} a^\top q \quad (23)$$

$$\frac{1}{2} H_M(q) = \min_{1a^\top + a1^\top \succeq L} a^\top q = \min_{a \in \mathcal{P}_L^M} a^\top q, \quad (24)$$

where the second expression is given by Prop. A.3. To show that $2H_M \leq H_L$ we will show that $\mathcal{P}_L \subseteq \mathcal{P}_L^M$. If $a \in \mathcal{P}_L$, then there exists $p \in \Delta$ such that $a \succeq Lp$. Moreover, if L is a distance, it means that the triangle inequality $L(y, y') \leq L(y, z) + L(z, y')$, which is equivalent to $L \preceq 1L_z^\top + L_z 1^\top$ for all $z \in \mathcal{Z}$. This can also be written as

$$L \preceq 1p^\top L + Lp1^\top, \quad \forall p \in \Delta. \quad (25)$$

Finally, note that if $Lp \preceq a$, then $Lp1^\top \preceq a1^\top$, and the same holds for its transpose $1p^\top L \preceq 1a^\top$. Hence, we obtain that $L \preceq 1a^\top + a1^\top$, which is equivalent to $a \in \mathcal{P}_L^M$. □

Proposition B.12. *If L is a distance, then $\frac{1}{2}H_M(q) = H_L(q)$, for all $q \in \Delta$ such that $\|q\|_\infty \geq 1/2$. Moreover, under this condition on q , it is calibrated with the argmax decoding.*

Proof. Combining Lemma B.10 and Lemma B.11 gives

$$H_M(q) \leq 2H_L(q) = 2L_y^\top q,$$

for all $q \in \Delta$ such that $p_y \geq 1/2 \geq p_z$ for all $z \neq y$. Hence, in order to prove the equality at these dominant label points, we just need to find a matrix $Q \in U(q, q)$ such that $\langle L, Q \rangle_F = 2L_y^\top q$. We define the matrix Q as

$$Q_{ij} = \begin{cases} q_i & \text{if } (i = y) \wedge (i \neq j) \\ q_j & \text{if } (j = y) \wedge (j \neq i) \\ 2q_y - 1 & \text{if } i = j = y \\ 0 & \text{otherwise.} \end{cases}.$$

The matrix Q has q at the y -th row and y -th column and $2q_y - 1$ at the crossing point, instead of $2q_y$. The matrix is in $U(q, q)$ as the sum of the rows and columns gives q and it is non-negative because $2q_y - 1 \geq 0$ by assumption. Moreover, the objective satisfies

$$\begin{aligned} \langle L, Q \rangle_F &= \sum_{y \in \mathcal{Y}} L_{y'}^\top Q_{y'} = L_y^\top Q_y + \sum_{y' \neq y} L_{y'} Q_{y'} \\ &= L_y^\top q + \sum_{y' \neq y} q_{y'} L_{y'} e_y \\ &= L_y^\top q + \sum_{y' \neq y} q_{y'} L(y, y') = 2L_y^\top q. \end{aligned}$$

The first part of the result follows. For the second part, we show that $v_M^*(q) = -L_y$ at the points q satisfying the assumption. Note that we have $H_M(q) = 2L_y^\top q = S(v^*(q))^\top q$, so we only need to show $S_M(-L_y) = 2L_y$. For every $z \in \mathcal{Y}$, we have

$$\begin{aligned} S_M(-L_y, z) &= \max_{y' \in \mathcal{Y}} L(z, y') + (-L_y)^\top e_{y'} - (-L_y)^\top e_z \\ &= \max_{y' \in \mathcal{Y}} \{L(z, y') - L(y, y')\} + L(y, z) \\ &= 2L(y, z), \end{aligned}$$

where at the last step we have used $L(z, y') - L(y, y') \leq L(y, y')$ as L is a distance and so the maximization is achieved at $y' = y$. \square

C PROOFS ON RESTRICTED-MAX-MARGIN LOSS

The following assumption **A1** will be key to prove our consistency results.

Assumption A1: If q is an extreme point of $\Delta(y')$ for some $y' \in \mathcal{Y}$, then

$$\{q \in \Delta(y)\} \vee \{q_y = 0\}, \quad \forall y \in \mathcal{Y}. \quad (26)$$

The following Lemma C.1 will be useful for the results below.

Lemma C.1. *If **A1** is satisfied, then $\Delta(y) \cap \Delta(y') \neq \emptyset$ for all $y, y' \in \mathcal{Y}$.*

Proof. If $\Delta(y) \cap \Delta(y') = \emptyset$ and **A1** is satisfied, then for every q extreme point of $\Delta(y')$ we have that $q_y = 0$. Hence, the prediction set $\Delta(y)$ is included in the non full-dimensional polyhedron $\Delta \cap \{e_y = 0\}$. As $\Delta = \cup_{y \in \mathcal{Y}} \Delta(y)$, this implies that the point $e_{y'} \in \Delta(y')$ must be necessarily included in another $\Delta(z)$, which can only be possible if $L(z, y) = 0$. However, by assumption $L(z, y) = 0$ if and only if $z = y$. \square

The following Lemma C.2 shows that under **A1**, the Restricted-Max-Margin loss embeds the loss L , which in turn implies consistency.

Lemma C.2. *Assume **A1**. If q is an extreme point of $\Delta(y)$ for some $y \in \mathcal{Y}$, then*

$$qq^\top \in \arg \max_{Q \in U(q, q) \cap \mathcal{C}_L} \langle L, Q \rangle_{\text{F}} \quad \text{and} \quad \Omega_{\text{MM}}(q) = \Omega_{\text{RM}}(q).$$

Proof. The matrix qq^\top belongs to $U(q, q)$ as $qq^\top \mathbf{1} = q$ and $qq^\top \succeq 0$ and to \mathcal{C}_L by assumption. Let $z \in y^*(q)$. We have that

$$\begin{aligned} -\Omega_{\text{RM}}(q) &= \max_{Q \in U(q, q) \cap \mathcal{C}_L} \langle L, Q \rangle_{\text{F}} \\ &\geq \langle L, qq^\top \rangle_{\text{F}} = \sum_y q_y L_y^\top q = \sum_y q_y L_z^\top q \\ &= L_z^\top (\mathbf{1}^\top qq^\top) = L_z^\top q = -\Omega_{\text{MM}}(q) \end{aligned}$$

We have shown that $\Omega_{\text{RM}}(q) \leq \Omega_{\text{MM}}(q)$. Combining with Proposition 3.4 that states $\Omega_{\text{RM}} \geq \Omega_{\text{MM}}$, we obtain that $\Omega_{\text{RM}}(q) = \Omega_{\text{MM}}(q)$. \square

Theorem C.3. *If **A1** is satisfied, the Restricted-Max-Margin loss embeds L with embedding $\psi(y) = -L_y$ and the loss is consistent to L under the argmax decoding.*

Proof. We split the proof into two parts.

First part: S_{RM} embeds L . Let $z \in y^*(q)$, so that $q \in \Delta(z)$. We can write q as a convex combination of extreme points of the polytope $\Delta(z)$ as

$$q = \sum_{i=1}^m \alpha_i q_i,$$

where $\alpha \in \Delta_m$ and q_i is an extreme point of $\Delta(z)$. The matrix $Q = \sum_{i=1}^m \alpha_i q_i q_i^\top$ belongs to $U(q, q) \cap \mathcal{C}_L$ as:

- $Q \in U(q, q)$: We have $Q\mathbf{1} = \sum_{i=1}^m \alpha_i q_i q_i^\top \mathbf{1} = \sum_{i=1}^m \alpha_i q_i = q$, the same holds for Q^\top and $\sum_{i=1}^m \alpha_i q_i q_i^\top \succeq 0$.
- $Q \in \mathcal{C}_L$: For all $y \in \mathcal{Y}$, we have $(1L_y^\top - L)Q_y = \sum_{i=1}^m \alpha_i \underbrace{q_{i,y}(1L_y^\top - L)q_i}_{\preceq 0} \preceq 0$.

Moreover, we obtain:

$$\begin{aligned}
 -\Omega_{\text{RM}}(q) &= \max_{Q \in U(q, q) \cap \mathcal{C}_L} \langle L, Q \rangle_{\text{F}} \\
 &\geq \langle L, Q \rangle_{\text{F}} = \sum_{i=1}^m \alpha_i \langle L, q_i q_i^{\top} \rangle_{\text{F}} \\
 &= - \sum_{i=1}^m \alpha_i \Omega_{\text{RM}}(q_i) = - \sum_{i=1}^m \alpha_i \Omega_{\text{MM}}(q_i) \\
 &= \sum_{i=1}^m \alpha_i L_z^{\top} q_i = L_z^{\top} q = -\Omega_{\text{MM}}(q)
 \end{aligned}$$

We have shown $\Omega_{\text{RM}} \leq \Omega_{\text{MM}}$. Combining with Proposition 3.4 that states $\Omega_{\text{RM}} \geq \Omega_{\text{MM}}$, we obtain $\Omega_{\text{RM}} = \Omega_{\text{MM}}$.

Second part: the embedding is $\psi(\mathbf{y}) = -L_{\mathbf{y}}$. By Prop. A.1, we only need to show that $\psi(z) = -L_z$, i.e.,

$$S_{\text{RM}}(-L_z, y) = \sup_{q \in \Delta(y)} L_y^{\top} q + (-L_z)^{\top} q - (-L_z)^{\top} e_y = L(y, z),$$

which holds whenever

$$\max_{q \in \Delta(y)} (L_y - L_z)^{\top} q = 0, \quad (27)$$

for all $y, z \in \mathcal{Y}$. Note that by construction $(L_y - L_z)^{\top} q \leq 0$ for all $q \in \Delta(y)$. Moreover, by Lemma C.1 we have that $\Delta(z) \cap \Delta(y) \neq \emptyset$, so there exists $q \in \Delta(y)$ with $L_y^{\top} q = L_z^{\top} q$.

Finally, the argmax decoding is consistent as it is an inverse of the embedding $\psi(y) = -L_y$ as

$$d(\psi(y)) = \arg \max_{y' \in \mathcal{Y}} -L(y, y') = \arg \min_{y' \in \mathcal{Y}} L(y, y') = y.$$

□

Proposition C.4. *Assume that $q \in \Delta(y) \implies q_y > 0$ for all $q \in \Delta$. Then **A1** is satisfied.*

Proof. We will prove that if the Assumption is not satisfied then it exists a vertex of $\Delta(y)$ for some $y \in \mathcal{Y}$ such that $S \cap T \neq \emptyset$. If the Assumption is not satisfied at vertex q , then $\{q \notin \Delta(y)\} \wedge \{q_y > 0\}$, which means in particular that $S \cup T \subsetneq [k]$. This necessarily means that $S \cap T \neq \emptyset$ because we must have $|S| + |T| \geq k$ to have maximal rank as q is a vertex. □

Proposition C.5. *Consistency of the Max-Margin implies consistency of Restricted-Max-Margin.*

Proof. From Prop. B.1 and Prop. B.5, we know that if the Max-Margin loss is consistent to L , then the extreme points of the prediction sets $\Delta(y)$'s have to be of the form $1/2(e_y + e_{y'})$. We will see that in this case **(A1)** is always satisfied. Indeed, if q is an extreme point of a prediction set, then is of the form $q = 1/2(e_y + e_{y'})$, which satisfies $\{q \in \Delta(z)\} \vee \{q_z = 0\}$ for all $z \in \mathcal{Y}$, because $q \in \Delta(z)$ if $z \in \{y, y'\}$ and $q_z = 0$ otherwise. □