
Quantifying Learning Guarantees for Convex but Inconsistent Surrogates

Kirill Struminsky
NRU HSE,^{*} Moscow, Russia

Simon Lacoste-Julien[†]
MILA and DIRO
Université de Montréal, Canada

Anton Osokin
NRU HSE,^{*‡} Moscow, Russia
Skoltech,[§] Moscow, Russia

Abstract

We study consistency properties of machine learning methods based on minimizing convex surrogates. We extend the recent framework of Osokin et al. [14] for the quantitative analysis of consistency properties to the case of inconsistent surrogates. Our key technical contribution consists in a new lower bound on the calibration function for the quadratic surrogate, which is non-trivial (not always zero) for inconsistent cases. The new bound allows to quantify the level of inconsistency of the setting and shows how learning with inconsistent surrogates can have guarantees on sample complexity and optimization difficulty. We apply our theory to two concrete cases: multi-class classification with the tree-structured loss and ranking with the mean average precision loss. The results show the approximation-computation trade-offs caused by inconsistent surrogates and their potential benefits.

1 Introduction

Consistency is a desirable property of any statistical estimator, which informally means that in the limit of infinite data, the estimator converges to the correct quantity. In the context of machine learning algorithms based on surrogate loss minimization, we usually use the notion of Fisher consistency, which means that the exact minimization of the expected surrogate loss leads to the exact minimization of the actual task loss. It can be shown that Fisher consistency is closely related to the question of infinite-sample consistency (a.k.a. classification calibration) of the surrogate loss with respect to the task loss (see [2, 17] for a detailed review).

The property of infinite-sample consistency (which we will refer to as simply consistency) shows that the minimization of a particular surrogate is the right problem to solve, but it becomes especially attractive when one can actually minimize the surrogate, which is the case, e.g, when the surrogate is convex. Consistency of convex surrogates has been the central question of many studies for such problems as binary classification [2, 24, 19], multi-class classification [23, 21, 1, 17], ranking [11, 4, 5, 18, 15] and, more recently, structured prediction [7, 14].

Recently, Osokin et al. [14] have pinpointed that in some cases minimizing a consistent convex surrogate might be not sufficient for efficient learning. In particular, when the number of possible predictions is large (which is typically the case in the settings of structured prediction and ranking) reaching adequately small value of the expected task loss can be practically impossible, because one would need to optimize the surrogate to high accuracy, which requires an intractable number of iterations of the optimization algorithm.

^{*}National Research University Higher School of Economics

[†]CIFAR Fellow

[‡]Samsung-HSE Joint Lab

[§]Skolkovo Institute of Science and Technology

It also turns out [14] that the possibility of efficient learning is related to the structure of the task loss. The 0-1 loss, which does not make distinction between different kinds of errors, shows the worst case behavior. However, more structured losses, e.g., the Hamming distance between sequence labelings, allow efficient learning if the score vector is designed appropriately (for the Hamming distance, the score for a complete configuration should be decomposable into the sum of scores for individual elements).

However, the analysis of Osokin et al. [14] gives non-trivial conclusions only for consistent surrogates. At the same time it is known that inconsistent surrogates often work well in practice (for example, the Crammer-Singer formulation of multi-class SVM [8], or its generalization structured SVM [20, 22]). There have indeed been several works to analyze inconsistent surrogates [12, 18, 5, 14], but they usually end the story with proving that some surrogate (or a family of surrogates) is not consistent.

Contributions. In this work, we look at the problem from a more quantitative angle and analyze to which extent inconsistent surrogates can be useful for learning. We focus on the same setting as [14] and generalize their results to the case of inconsistent surrogates (their bounds are trivial for these cases) to be able to draw non-trivial conclusions. The main technical contribution consists in a tighter lower bound on the calibration function (Theorem 3), which is strictly more general than the bound of [14]. Notably, our bound is non-trivial in the case when the surrogate is not consistent and quantifies to which degree learning with inconsistent surrogates is possible. We further study the behavior of our bound in two practical scenarios: multi-class classification with a tree-structured loss and ranking with the mean average precision (mAP) loss. For the tree-structured loss, our bound shows that there can be a trade-off between the best achievable accuracy and the speed of convergence. For the mAP loss, we use our tools to study the (non-)existence of consistent convex surrogates of a particular dimension (an important issue for the task of ranking [11, 4, 5, 18, 17]) and quantify to which extent our quadratic surrogate with the score vector of insufficient dimension is consistent.

This paper is organized as follows. First, we introduce the setting we work with in Section 2 and review the key results of [14] in Section 3. In Section 4, we prove our main theoretical result, which is a new lower bound on the calibration function. In Section 5, we analyze the behavior of our bound for the two different settings: multi-class classification and ranking (the mean average precision loss). Finally, we review the related works and conclude in Section 6.

2 Notation and Preliminaries

In this section, we introduce our setting, which closely follows [14]. We denote the input features by $x \in \mathcal{X}$ where \mathcal{X} is the input domain. The particular structure of \mathcal{X} is not of the key importance for this study. The output variables, that are in the center of our analysis, will be denoted by $\hat{y} \in \hat{\mathcal{Y}}$ with $\hat{\mathcal{Y}}$ being the set of possible predictions or the output domain.⁵ In such settings as structured prediction or ranking, the predictions are very high-dimensional and with some structure that is useful to model explicitly (for example, a sequence, permutation or image).

The central object of our study is the *loss function* $L(\hat{y}, y) \geq 0$ that represents the cost of making the prediction $\hat{y} \in \hat{\mathcal{Y}}$ when the ground-truth label is $y \in \mathcal{Y}$. Note that in some applications of interest the sets $\hat{\mathcal{Y}}$ and \mathcal{Y} are different. For example, in ranking with the mean average precision (mAP) loss function (see Section 5.2 and, e.g., [18] for the details), the set $\hat{\mathcal{Y}}$ consists of all the permutations of the items (to represent the ranking itself), but the set \mathcal{Y} consists of all the subsets of items (to represent the set of relevant items, which is the ground-truth annotation in this setting). In this paper, we only study the case when both $\hat{\mathcal{Y}}$ and \mathcal{Y} are finite. We denote the cardinality of $\hat{\mathcal{Y}}$ by k , and the cardinality of \mathcal{Y} by m . In this case, the loss function can be encoded as a matrix L of size $k \times m$.

In many applications of interest, both quantities k and m are exponentially large in the size of the natural dimension of the input x . For example, in the task of sequence labeling, both k and m are equal to the number of all possible sequences of symbols from a finite alphabet. In the task of ranking (the mAP formulation), k is equal to the number of permutations of items and m is equal to the number of item subsets.

⁵The output domain $\hat{\mathcal{Y}}$ itself can depend on the vector of input features x (for example, if x can represent sequences of different lengths and the length of the output sequence has to equal the length of the input), but we will not use this dependency and omit it for brevity.

Following usual practices, we work with the prediction model defined by a (learned) vector-valued *score function* $\mathbf{f} : \mathcal{X} \rightarrow \mathbb{R}^k$, which defines a scalar score $\mathbf{f}_{\hat{\mathbf{y}}}(\mathbf{x})$ for each possible output $\hat{\mathbf{y}} \in \hat{\mathcal{Y}}$. The final prediction is then chosen as an output configuration with the maximal score:

$$\text{pred}(\mathbf{f}(\mathbf{x})) := \underset{\hat{\mathbf{y}} \in \hat{\mathcal{Y}}}{\operatorname{argmax}} \mathbf{f}_{\hat{\mathbf{y}}}(\mathbf{x}). \quad (1)$$

If the maximal score is given by multiple outputs $\hat{\mathbf{y}}$ (so-called *ties*), the predictor follows a simple deterministic tie-breaking rule and picks the output appearing first in some predefined ordering on $\hat{\mathcal{Y}}$.

In this setup, learning consists in finding a *score function* \mathbf{f} for which the predictor gives the smallest expected loss with features \mathbf{x} and labels \mathbf{y} coming from an unknown data-generating distribution \mathcal{D} :

$$\mathcal{R}_L(\mathbf{f}) := \mathbb{E}_{(\mathbf{x}, \mathbf{y}) \sim \mathcal{D}} L(\text{pred}(\mathbf{f}(\mathbf{x})), \mathbf{y}). \quad (2)$$

The quantity $\mathcal{R}_L(\mathbf{f})$ is usually referred to as the actual (or population) *risk* based on the loss L . Minimizing the actual risk directly is usually difficult (because of non-convexity and non-continuity of the predictor (1)). The standard approach is to substitute (2) with another objective, a *surrogate risk* (or the Φ -risk), which is easier for optimization (in this paper, we only consider convex surrogates):

$$\mathcal{R}_\Phi(\mathbf{f}) := \mathbb{E}_{(\mathbf{x}, \mathbf{y}) \sim \mathcal{D}} \Phi(\mathbf{f}(\mathbf{x}), \mathbf{y}), \quad (3)$$

where we will refer to the function $\Phi : \mathbb{R}^k \times \mathcal{Y} \rightarrow \mathbb{R}$ as the *surrogate loss*. To make the minimization of (3) well-defined, we will always assume the surrogate loss Φ to be bounded from below and continuous.

The surrogate loss should be chosen in such a way that the minimization of (3) also leads to the minimization of (2), i.e., to the solution of the original problem. The property of consistency of the surrogate loss is an approach to formalize this intuition, i.e., to guarantee that no matter the data-generating distribution, minimizing (3) w.r.t. \mathbf{f} implies minimizing (2) w.r.t. \mathbf{f} as well (both of these are possible only in the limit of infinite data and computational budget). Osokin et al. [14] quantified what happens if the surrogate risk is minimized approximately by translating the optimization error of (3) to the optimization error of (2). The main goal of this paper is to generalize this analysis to the cases when the surrogate is not consistent and to show that there can be trade-offs between the minimum value of the actual risk that can be achieved by minimizing an inconsistent surrogate and the speed with which this minimum can be achieved.

3 Calibration Functions and Consistency

In this section, we review the approach of Osokin et al. [14] for studying consistency in the context of structured prediction. The first part of the analysis establishes the connection between the minimization of the actual risk \mathcal{R}_L (2) and the surrogate risk \mathcal{R}_Φ (3) via the so-called *calibration function* (see Definition 1 [14, and references therein]). This step is usually called *non-parametric* (or pointwise) because it does not explicitly model the dependency of the scores $\mathbf{f} := \mathbf{f}(\mathbf{x})$ on the input variables \mathbf{x} . The second part of the analysis establishes the connection with an optimization algorithm allowing to make a statement about how many iterations would be enough to find a predictor that is (in expectation) within ε of the global minimum of the actual risk \mathcal{R}_L .

Non-parametric analysis. The standard non-parametric setting considers all measurable score functions \mathbf{f} to effectively ignore the dependency on the features \mathbf{x} . As noted by [14], it is beneficial to consider a restricted set of the score functions $\mathfrak{F}_{\mathcal{F}}$ that consists of all vector-valued Borel measurable functions $\mathbf{f} : \mathcal{X} \rightarrow \mathcal{F}$ where $\mathcal{F} \subseteq \mathbb{R}^k$ is a subspace of allowed score vectors. Compatibility of the subspace \mathcal{F} and the loss function L will be a crucial point of this paper. Note that the analysis is still non-parametric because the dependence on \mathbf{x} is not explicitly modeled.

Within the analysis, we will use the *conditional* actual and surrogate risks defined as the expectations of the corresponding losses w.r.t. a categorical distribution \mathbf{q} on the set of annotations \mathcal{Y} , $m := |\mathcal{Y}|$:

$$\ell(\mathbf{f}, \mathbf{q}) := \sum_{\mathbf{y}=1}^m q_{\mathbf{y}} L(\text{pred}(\mathbf{f}), \mathbf{y}), \quad \phi(\mathbf{f}, \mathbf{q}) := \sum_{\mathbf{y}=1}^m q_{\mathbf{y}} \Phi(\mathbf{f}, \mathbf{y}). \quad (4)$$

Hereinafter, we represent an m -dimensional categorical distribution \mathbf{q} as a point in the probability simplex Δ_m and use the symbol $q_{\mathbf{y}}$ to denote the probability of the \mathbf{y} -th outcome. Using this notation, we can rewrite the risk \mathcal{R}_L and surrogate risk \mathcal{R}_Φ as

$$\mathcal{R}_L(\mathbf{f}) = \mathbb{E}_{\mathbf{x} \sim \mathcal{D}_{\mathcal{X}}} \ell(\mathbf{f}(\mathbf{x}), \mathbf{P}_{\mathcal{D}}(\cdot | \mathbf{x})), \quad \mathcal{R}_\Phi(\mathbf{f}) = \mathbb{E}_{\mathbf{x} \sim \mathcal{D}_{\mathcal{X}}} \phi(\mathbf{f}(\mathbf{x}), \mathbf{P}_{\mathcal{D}}(\cdot | \mathbf{x})), \quad (5)$$

where \mathcal{D}_X is the marginal distribution of x and $\mathbf{P}_{\mathcal{D}}(\cdot | x)$ denotes the conditional distribution of y given x (both defined for the joint data-generating distribution \mathcal{D}).

For each score vector $\mathbf{f} \in \mathcal{F}$ and a distribution $\mathbf{q} \in \Delta_m$ over ground-truth labels, we now define the *excess* actual and surrogate risks

$$\delta\phi(\mathbf{f}, \mathbf{q}) = \phi(\mathbf{f}, \mathbf{q}) - \inf_{\hat{\mathbf{f}} \in \mathcal{F}} \phi(\hat{\mathbf{f}}, \mathbf{q}), \quad \delta\ell(\mathbf{f}, \mathbf{q}) = \ell(\mathbf{f}, \mathbf{q}) - \inf_{\hat{\mathbf{f}} \in \mathbb{R}^k} \ell(\hat{\mathbf{f}}, \mathbf{q}), \quad (6)$$

which show how close the current conditional actual and surrogate risks are to the corresponding minimal achievable conditional risks (depending only on the distribution \mathbf{q}). Note that the two infima in (6) are defined w.r.t. different sets of score vectors. For the surrogate risk, the infimum is taken w.r.t. the set of allowed scores \mathcal{F} capturing only the scores obtainable by the learning process. For the actual risk, the infimum is taken w.r.t. the set of all possible scores \mathbb{R}^k including score vectors that cannot be learned. This distinction is important when analyzing inconsistent surrogates and allows to characterize the *approximation error* of the selected function class.⁶

We are now ready to define the *calibration function*, which is the final object of the non-parametric part of the analysis. Calibration functions directly show how well one needs to minimize the surrogate risk to guarantee that the excess of the actual risk is smaller than ε .

Definition 1 (Calibration function, [14]). *For a task loss L , a surrogate loss Φ , a set of feasible scores \mathcal{F} , the calibration function $H_{\Phi, L, \mathcal{F}}(\varepsilon)$ is defined as:*

$$H_{\Phi, L, \mathcal{F}}(\varepsilon) := \inf_{\mathbf{f} \in \mathcal{F}, \mathbf{q} \in \Delta_m} \delta\phi(\mathbf{f}, \mathbf{q}) \quad (7)$$

$$\text{s.t.} \quad \delta\ell(\mathbf{f}, \mathbf{q}) \geq \varepsilon, \quad (8)$$

where $\varepsilon \geq 0$ is the target accuracy. We set $H_{\Phi, L, \mathcal{F}}(\varepsilon)$ to $+\infty$ when the feasible set is empty.

By construction, $H_{\Phi, L, \mathcal{F}}$ is non-decreasing on $[0, +\infty)$, $H_{\Phi, L, \mathcal{F}}(\varepsilon) \geq 0$ and $H_{\Phi, L, \mathcal{F}}(0) = 0$. The calibration function also provides the so-called *excess risk bound*

$$H_{\Phi, L, \mathcal{F}}(\delta\ell(\mathbf{f}, \mathbf{q})) \leq \delta\phi(\mathbf{f}, \mathbf{q}), \quad \forall \mathbf{f} \in \mathcal{F}, \quad \forall \mathbf{q} \in \Delta_m, \quad (9)$$

which implies the formal connection between the surrogate and task risks [14, Theorem 2].

The calibration function can fully characterize consistency of the setting defined by the surrogate loss, the subspace of scores and the task loss. The maximal value of ε at which the calibration function $H_{\Phi, L, \mathcal{F}}(\varepsilon)$ equals zero shows the best accuracy on the actual loss that can be obtained [14, Theorem 6]. The notion of level- η consistency captures this effect.

Definition 2 (level- η consistency, [14]). *A surrogate loss Φ is consistent up to level $\eta \geq 0$ w.r.t. a task loss L and a set of scores \mathcal{F} if and only if the calibration function satisfies $H_{\Phi, L, \mathcal{F}}(\varepsilon) > 0$ for all $\varepsilon > \eta$ and there exists $\hat{\varepsilon} > \eta$ such that $H_{\Phi, L, \mathcal{F}}(\hat{\varepsilon})$ is finite.*

The case of level-0 consistency corresponds to the classical consistent surrogate and Fisher consistency. When $\eta > 0$, the surrogate is not consistent, meaning that the actual risk cannot be minimized globally. However, Osokin et al. [14, Appendix E.4] give an example where even though constructing a consistent setting is possible (by the choice of the score subspace \mathcal{F}), it might still be beneficial to use only a level- η consistent setting because of the exponentially faster growth of the calibration function. The main contribution of this paper is a lower bound on the calibration function (Theorem 3), which is non-zero for $\eta > 0$ and thus can be used to obtain convergence rates in inconsistent settings.

Optimization and learning guarantees; normalizing the calibration function. Osokin et al. [14] note that the scale of the calibration function is not defined, i.e., if one multiplies the surrogate loss by some positive constant, the calibration function is multiplied by the same constant as well. One way to define a “natural normalization” is to use a scale-invariant convergence rate of a stochastic optimization algorithm. Osokin et al. [14, Section 3.3] applied the classical online ASGD [13] (under the well-specification assumption) and got the sample complexity (and the convergence rate of ASGD at the same time) result saying that N^* steps of ASGD are sufficient to get ε -accuracy on the *task loss* (in expectation), where N^* is computed as follows:

$$N^* := \frac{4D^2M^2}{\tilde{H}_{\Phi, L, \mathcal{F}}^2(\varepsilon)}. \quad (10)$$

⁶Note that Osokin et al. [14] define the excess risks by taking both infima w.r.t. the set of allowed scores \mathcal{F} , which is subtly different from us. The results of the two setups are equivalent in the cases of consistent surrogates, which are the main focus of Osokin et al. [14], but can be different in inconsistent cases.

Here the quantity N^* depends on a convex lower bound $\tilde{H}_{\Phi, L, \mathcal{F}}(\varepsilon)$ on the calibration function $H_{\Phi, L, \mathcal{F}}(\varepsilon)$ and the constants D, M , which appear in the convergence rate of ASGD: D is an upper bound on the norm of an optimal solution and M^2 is an upper bound on the expected square norm of the stochastic gradient. Osokin et al. [14] show how to bound the constant DM for a very specific quadratic surrogate defined below (see Section 3.1).

3.1 Bounds for the Quadratic Surrogate

The major complication in applying and interpreting the theoretical results presented in Section 3 is the complexity of computing the calibration function. Osokin et al. [14] analyzed the calibration function only for the quadratic surrogate

$$\Phi_{\text{quad}}(\mathbf{f}, \hat{\mathbf{y}}) := \frac{1}{2k} \|\mathbf{f} + L(\cdot, \mathbf{y})\|_2^2 = \frac{1}{2k} \sum_{\hat{\mathbf{y}} \in \hat{\mathcal{Y}}} (f_{\hat{\mathbf{y}}}^2 + 2f_{\hat{\mathbf{y}}}L(\hat{\mathbf{y}}, \mathbf{y}) + L(\hat{\mathbf{y}}, \mathbf{y})^2). \quad (11)$$

For any task loss L , this surrogate is consistent whenever the subspace of allowed scores is rich enough, i.e., the subspace of scores \mathcal{F} fully contains $\text{span}(L)$. To connect with optimization, we assume a parametrization of the subspace \mathcal{F} as a span of the columns of some matrix F , i.e., $\mathcal{F} = \text{span}(F) = \{\mathbf{f} = F\boldsymbol{\theta} \mid \boldsymbol{\theta} \in \mathbb{R}^r\}$.⁷ In the interesting settings, the dimension r is much smaller than both k and m . Note that to compute the gradient of the objective (11) w.r.t. the parameters $\boldsymbol{\theta}$, one needs to compute matrix products $F^T F \in \mathbb{R}^{r \times r}$ and $F^T L(\cdot, \mathbf{y}) \in \mathbb{R}^r$, which are usually both of feasible sizes, but require exponentially big sum (k summands) inside. Computing these quantities can be seen as some form of inference required to run the learning process.

Osokin et al. [14] proved a lower bound on the calibration functions for the quadratic surrogate (11) [14, Theorem 7], which we now present to contrast our result presented in Section 4. When the subspace of scores \mathcal{F} contains $\text{span}(L)$, $\text{span}(L) \subseteq \mathcal{F}$, implying that the setting is consistent, the calibration function is bounded from below by $\min_{i \neq j} \frac{\varepsilon^2}{2k \|P_{\mathcal{F}} \Delta_{ij}\|_2^2}$, where $P_{\mathcal{F}}$ is the orthogonal projection on the subspace \mathcal{F} and $\Delta_{ij} := \mathbf{e}_i - \mathbf{e}_j \in \mathbb{R}^k$ with \mathbf{e}_c being the c -th basis vector of the standard basis in \mathbb{R}^k . They also showed that for some very structured losses (Hamming and block 0-1 losses), the quantity $k \|P_{\mathcal{F}} \Delta_{ij}\|_2^2$ is not exponentially large and thus the calibration function suggests that efficient learning is possible. One interesting case not studied by Osokin et al. [14] is the situation where the subspace of scores \mathcal{F} does not fully contain the subspace $\text{span}(L)$. In this case, the surrogate might not be consistent but still lead to effective and efficient practical algorithms.

Normalizing the calibration function. The normalization constant DM appearing in (10) can also be computed for the quadratic surrogate (11) under the assumption of well-specification (see [14, Appendix F] for details). In particular, we have $DM = L_{\max}^2 \xi(\kappa(F) \sqrt{r} R Q_{\max})$, $\xi(z) = z^2 + z$, where L_{\max} denotes the maximal value of all elements in L , $\kappa(F)$ is the condition number of the matrix F and r in an upper bound on the rank of \mathcal{F} . The constants R and Q_{\max} come from the kernel ASGD setup and, importantly, depend only on the data distribution, but not on the loss L or score matrix F . Note that for a given subspace \mathcal{F} , the choice of matrix F is arbitrary and it can always be chosen as an orthonormal basis of \mathcal{F} giving a $\kappa(F)$ of one. However, such F can lead to inefficient prediction (1), which makes the whole framework less appealing. Another important observation coming from the value of DM is the justification of the $\frac{1}{k}$ scaling in front of the surrogate (11).

4 Calibration Function for Inconsistent Surrogates

Our main result generalizes the Theorem 7 of [14] to the case of inconsistent surrogates (the key difference consists in the absence of the assumption $\text{span}(L) \subseteq \mathcal{F}$).

Theorem 3 (Lower bound on the calibration function $H_{\Phi_{\text{quad}}, L, \mathcal{F}}(\varepsilon)$). *For any task loss L , its quadratic surrogate Φ_{quad} , and a score subspace \mathcal{F} , the calibration function is bounded from below:*

$$H_{\Phi_{\text{quad}}, L, \mathcal{F}}(\varepsilon) \geq \min_{i \neq j} \max_{v \geq 0} \frac{(\varepsilon v - \xi_{ij}(v))_+^2}{2k \|P_{\mathcal{F}} \Delta_{ij}\|_2^2}, \quad \text{where} \quad \xi_{ij}(v) := \|L^T(v \mathbf{I}_k - P_{\mathcal{F}}) \Delta_{ij}\|_{\infty}, \quad (12)$$

where $P_{\mathcal{F}}$ is the orthogonal projection on the subspace \mathcal{F} , $(x)_+^2 := [x > 0]x^2$ is the truncation of the parabola to its right branch and $\Delta_{ij} := \mathbf{e}_i - \mathbf{e}_j \in \mathbb{R}^k$ with $\mathbf{e}_c \in \mathbb{R}^k$ being the c -th column of the

⁷We do a pointwise analysis in this section, so we are not modeling the dependence of $\boldsymbol{\theta}$ on the features \mathbf{x} . However, in an actual implementation, the vector $\boldsymbol{\theta}$ should be a function of the features \mathbf{x} coming from some flexible family such as a RKHS or some neural networks.

identity matrix \mathbf{I}_k . By convention, if both numerator and denominator of (12) equal zero the whole bound equals zero. If only the denominator equals zero then the whole bound equals infinity (the particular pair of i and j is effectively not considered).

The proof of Theorem 3 starts with using the idea of [14] to compute the calibration function by solving a collection of convex quadratic programs (QPs). Then we diverge from the proof of [14] (because it leads to a non-informative bound in inconsistent settings). For each of the formulated QPs, we construct a dual by using the approach of Dorn [10]. The dual of Dorn is convenient for our needs because it does not require inverting the matrix defining the quadratic terms (compared to the standard Lagrangian dual). The complete proof is given in Appendix B.

Remark 4. The numerator of the bound (12) explicitly specifies the point at which the bound becomes non-zero, implying level- η consistency with $\eta = \frac{\xi_{ij}(v)}{v}$ for the values of i, j, v that are active for a particular ε . The quantity $\frac{v^2}{2k\|\mathbf{P}_{\mathcal{F}}\Delta_{ij}\|_2^2}$ bounds the weight of the ε^2 term in the calibration function after it leaves zero. Moving the quantity v defines the trade-off between the slope, which is related to the convergence speed of the algorithm, and the value of η defining the best achievable accuracy.

Remark 5. If we have conditions of Theorem 7 of [14] satisfied, i.e., $\text{span}(L) \subseteq \mathcal{F}$, then the vector $L^\top(\mathbf{I}_k - \mathbf{P}_{\mathcal{F}})\Delta_{ij}$ equals zero and $\xi_{ij}(v)$ becomes $|v - 1| \|\mathbf{L}^\top\Delta_{ij}\|_\infty$, which equals zero when $v = 1$. It might seem that having $v > 1$ can potentially give us a tighter lower bound than Theorem 7 [14] even in consistent cases. However, the quantity $\|\mathbf{L}^\top\Delta_{ij}\|_\infty$ upper bounds the maximal possible (w.r.t. the conditional distribution $\mathbf{P}_{\mathcal{D}}(\cdot | \mathbf{x})$) value of the excess task loss for a fixed pair i, j leading to the identity $v\varepsilon - |v - 1| \|\mathbf{L}^\top\Delta_{ij}\|_\infty = \|\mathbf{L}^\top\Delta_{ij}\|_\infty$ for $\varepsilon = \|\mathbf{L}^\top\Delta_{ij}\|_\infty$ and $v \geq 1$. Together with the convexity of the function $(x)_+^2$, this implies that the best possible value of v in consistent settings equals one.

Remark 6. Setting v in (12) to any non-negative constant gives a valid lower bound. In particular, setting v to 1 (while potentially making the bound less tight) highlights the separation between the weight of the quadratic term and the best achievable accuracy η . The bound now reads as follows:

$$H_{\Phi_{\text{quad}}, L, \mathcal{F}}(\varepsilon) \geq \min_{i \neq j} \frac{(\varepsilon - \xi_{ij})_+^2}{2k\|\mathbf{P}_{\mathcal{F}}\Delta_{ij}\|_2^2}, \quad \text{where} \quad \xi_{ij} := \|\mathbf{L}^\top(\mathbf{I}_k - \mathbf{P}_{\mathcal{F}})\Delta_{ij}\|_\infty. \quad (13)$$

Note that the weight of the ε^2 term now equals the corresponding coefficient of the bound of Theorem 7 [14]. Notably, this weight depends only on the score subspace \mathcal{F} , but not on the loss L .

5 Bounds for Particular Losses

5.1 Multi-Class Classification with the Tree-Structured Loss

As an illustration of the obtained lower bound (12), we consider the task of multi-class classification and the *tree-structured loss*, which is defined for a weighted tree built on labels (such trees on labels often appear in settings with large number of labels, e.g., extreme classification [6]). Leaves in the tree correspond to the class labels $\hat{y} \in \hat{\mathcal{Y}} = \mathcal{Y}$ and the loss function is defined as the length of the path ρ between the leaves, i.e., $L_{\text{tree}}(y, \hat{y}) := \rho(y, \hat{y})$. To compute the lower bound exactly, we assume that the number of children d_s and the weights of the edges connecting a node with its children $\frac{\alpha_s}{2}$ are equal for all the nodes of the same depth level $s = 0, \dots, D - 1$ (see Figure 2 in Appendix C for an example of such a tree) and that $\sum_{s=0}^{D-1} \alpha_s = 1$, which normalizes L_{max} to one.

To define the score matrix $\mathcal{F}_{\text{tree}, s_0}$, we set the *consistency depth* $s_0 \in \{1, \dots, D\}$ and restrict the scores \mathbf{f} to be equal for the groups (blocks) of leaves that have the same ancestor on the level s_0 . Let $B(i)$ be the set of leaves that have the same ancestor as a leaf i at the depth s_0 . With this notation, we have $\mathcal{F}_{\text{tree}, s_0} = \text{span} \{ \sum_{i \in B(j)} \mathbf{e}_i \mid j = 1, \dots, k \}$. Theorem 3 gives us the bound (see Appendix C):

$$H_{\Phi_{\text{quad}}, L_{\text{tree}}, \mathcal{F}_{\text{tree}, s_0}}(\varepsilon) \geq [\varepsilon > \eta_{s_0}] \frac{(\eta_{s_0} - \bar{\rho}_{s_0} + \alpha_{s_0-1})^2}{(\frac{\eta_{s_0}}{2} + \alpha_{s_0-1})^2} \frac{(\varepsilon - \frac{\eta_{s_0}}{2})_+^2}{4b_{s_0}}, \quad (14)$$

where $b_{s_0}, \bar{\rho}_{s_0} := \frac{1}{|B(j)|} \sum_{i \in B(j)} \rho(i, j) = \sum_{s=s_0}^{D-1} \alpha_s \frac{(\prod_{s'=s_0}^s d_{s'}) - 1}{\prod_{s'=s_0}^s d_{s'}}$ and $\eta_{s_0} := \max_{i \in B(j)} \rho(i, j) = \sum_{s=s_0}^{D-1} \alpha_s$ are the number of blocks, the average and maximal distance within a block, respectively.

Now we discuss the behavior of the bound (14) when changing the truncation level s_0 . With the growth of s_0 , the level of consistency η_{s_0} goes to 0 indicating that more labels can be distinguished. At the same time, we have $\frac{\eta_{s_0}}{2} \leq \bar{\rho}_{s_0}$ for the trees we consider and thus the coefficient in front of the

ε^2 term can be bounded from above by $\frac{1}{4bs_0}$, which means that the lower bound on the calibration function decreases at an exponential rate with the growth of s_0 . These arguments show the trade-off between the level of consistency and the coefficient of ε^2 in the calibration function.

Finally, note that the mixture of 0-1 and block 0-1 losses considered in [14, Appendix E.4] is an instance of the tree-structured loss with $D = 2$. Their bound [14, Proposition 17] matches (14) up to the difference in the definition of the calibration function (they do not have the $[\varepsilon > \eta_{s_0}]$ multiplier because they do not consider pairs of labels that fall in the same block).

5.2 Mean Average Precision (mAP) Loss for Ranking

The mAP loss, which is a popular way of measuring the quality of ranking, has attracted significant attention from the consistency point of view [4, 5, 18]. In the mAP setting, the ground-truth labels are binary vectors $\mathbf{y} \in \mathcal{Y} = \{0, 1\}^r$ that indicate the items relevant for the query (a subset of r items-to-rank) and the prediction consists in producing a permutation of items $\sigma \in \hat{\mathcal{Y}}$, $\hat{\mathcal{Y}} = \mathcal{S}_r$. The mAP loss is based on averaging the precision at different levels of recall and is defined as follows:

$$L_{\text{mAP}}(\sigma, \mathbf{y}) := 1 - \frac{1}{|\mathbf{y}|} \sum_{p: y_p=1}^r \frac{1}{\sigma(p)} \sum_{q=1}^{\sigma(p)} y_{\sigma^{-1}(q)} = 1 - \sum_{p=1}^r \sum_{q=1}^p \frac{1}{\max(\sigma(p), \sigma(q))} \frac{y_p y_q}{|\mathbf{y}|}, \quad (15)$$

where $\sigma(p)$ is the position of an item p in a permutation σ , σ^{-1} is the inverse permutation and $|\mathbf{y}| := \sum_{p=1}^r y_p$. The second identity provides a convenient form of writing the mAP loss [18]

showing that the loss matrix L_{mAP} is of rank at most $\frac{1}{2}r(r+1)$.⁸ The matrix $F_{\text{mAP}} \in \mathbb{R}^{r! \times \frac{1}{2}r(r+1)}$ such that $(F_{\text{mAP}})_{\sigma, pq} := \frac{1}{\max(\sigma(p), \sigma(q))}$ is a natural candidate to define the score subspace \mathcal{F} to get the consistent setting with the quadratic surrogate (11) (Eq. (15) implies that $\text{span}(L_{\text{mAP}}) = \text{span}(F_{\text{mAP}})$).

However, as noted in Section 6 of [18], although the matrix F_{mAP} is convenient from the consistency point of view (in the setup of [18]), it leads to the prediction problem $\max_{\sigma \in \mathcal{S}_r} (F_{\text{mAP}} \boldsymbol{\theta})_{\sigma}$, which is a quadratic assignment problem (QAP), and most QAPs are NP-hard.

To be able to predict efficiently, it would be beneficial to have the matrix F with r columns such that sorting the r -dimensional $\boldsymbol{\theta}$ would give the desired permutation. It appears that it is possible to construct such a matrix by selecting a subset of columns of matrix F_{mAP} . We define $F_{\text{sort}} \in \mathbb{R}^{r! \times r}$ by $(F_{\text{sort}})_{\sigma, p} := \frac{1}{\sigma(p)}$. A solution of the prediction problem $\max_{\sigma \in \mathcal{S}_r} (F_{\text{sort}} \boldsymbol{\theta})_{\sigma}$ is simply a permutation that sorts the elements of $\boldsymbol{\theta} \in \mathbb{R}^r$ in the decreasing order (this statement follows from the fact that we can always increase the score $(F_{\text{sort}} \boldsymbol{\theta})_{\sigma} = \sum_{p=1}^r \frac{\theta_p}{\sigma(p)}$ by swapping a pair of non-aligned items).

Most importantly for our study, the columns of the matrix F_{sort} are a subset of the columns of the matrix F_{mAP} , which indicates that learning with the convenient matrix F_{sort} might be sufficient for the mAP loss. In what follows, we study the calibration functions for the loss matrix L_{mAP} and score matrices F_{mAP} and F_{sort} . In Figure 1a-b, we plot the calibration functions for both F_{mAP} and F_{sort} and the lower bounds given by Theorem 3. All the curves were obtained for $r = 5$ (computing the exact values of the calibration functions is exponential in r).

Next, we study the behavior of the lower bound (12) for large values of r . In Lemma 13 of Appendix D, we show that the denominator of the bound (12) is not exponential in r (we have $2r! \|P_{\mathcal{F}_{\text{sort}}} \Delta_{\pi\omega}\|_2^2 = O(r)$). We also know that $\|P_{\mathcal{F}_{\text{sort}}} \Delta_{\pi\omega}\|_2^2 \leq \|P_{\mathcal{F}_{\text{mAP}}} \Delta_{\pi\omega}\|_2^2$ (because $\mathcal{F}_{\text{sort}}$ is a subspace of \mathcal{F}_{mAP}), which implies that the calibration function of the consistent setting grows not faster than the one of the inconsistent setting. We can also numerically compute a lower bound on the point η until which the calibration function is guaranteed to be zero (for this we simply pick two permutations π, ω and a labeling \mathbf{y} that delivers large values of $(L_{\text{mAP}}^T (\mathbf{I}_k - P_{\mathcal{F}_{\text{sort}}}) \Delta_{ij})_{\mathbf{y}} \leq \xi_{\pi, \omega}(1)$). Figure 1c shows that the level of inconsistency η grows with the growth of r , which makes the method less appealing for large-scale settings.

Finally, note that to run the ASGD algorithm for the quadratic surrogate (11), mAP loss and score matrix F_{sort} , we need to efficiently compute $F_{\text{sort}}^T F_{\text{sort}}$ and $F_{\text{sort}}^T L_{\text{mAP}}(:, \mathbf{y})$. Lemmas 11 and 12 (see Appendix D) provide linear in r time algorithms for doing this. The condition number of $\mathcal{F}_{\text{sort}}$ grows as $\Theta(\log r)$ keeping the sample complexity bound (10) well behaved.

⁸Ramaswamy & Agarwal [17, Proposition 21] showed that the rank of L_{mAP} is at least $\frac{1}{2}r(r+1) - 2$.

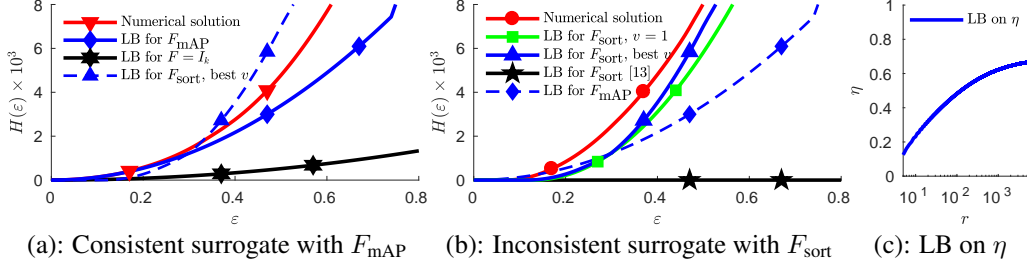


Figure 1: **Plot (a)** shows the calibration function $H_{\Phi_{\text{quad}}, L_{\text{mAP}}, \mathcal{F}_{\text{mAP}}}(\varepsilon)$ for L_{mAP} (red line) obtained numerically. The solid blue line [14, Theorem 7] is its lower bound, LB, and the solid black line is the worst case bound obtained for $F = I_r$ (which means not constructing an appropriate low-dimension \mathcal{F}). Difference between the blue and the black lines is exponential (proportional to $r!$). The dashed blue line illustrates the inconsistent surrogate (note that it is zero for small $\varepsilon > 0$, but then grows faster than the solid blue line – the consistent setting). **Plot (b)** shows the calibration function $H_{\Phi_{\text{quad}}, L_{\text{mAP}}, \mathcal{F}_{\text{sort}}}(\varepsilon)$ (red line) obtained numerically (this setting is level- η consistent for $\eta \approx 0.08$). The blue line (Theorem 3) is its lower bound for the optimal value of v and the green line is the bound for $v = 1$ (easier to obtain). The black line shows the zero-valued trivial bound from [14]. The dashed blue line shows $H_{\Phi_{\text{quad}}, L_{\text{mAP}}, \mathcal{F}_{\text{mAP}}}(\varepsilon)$ for the consistent surrogate to compare the two settings. Note that in both plots (a) and (b) the solid blue lines are the lower bounds of the corresponding calibration functions (red lines), but the dashed blue lines are not (shown for comparison purposes). **Plot (c)** shows a lower bound on the point η where the exact calibration function $H_{\Phi_{\text{quad}}, L_{\text{mAP}}, \mathcal{F}_{\text{sort}}}(\varepsilon)$ stops being zero, indicating the level of consistency (Definition 2).

6 Discussion

Related works. Despite a large number of works studying consistency and calibration in the context of machine learning, there have been relatively few attempts to obtain guarantees for inconsistent surrogates. The most popular approach is to study consistency under so-called *low noise* conditions. Such works show that under certain assumptions on the data generating distribution \mathcal{D} (usually these assumptions are on the conditional distribution of labels and are impossible to verify for real data) the surrogate of interest becomes consistent, whereas being inconsistent for general \mathcal{D} . Duchi et al. [11] established such a result for the value-regularized linear surrogate for ranking (which resembles the pairwise disagreement, PD, loss). Ramaswamy et al. [18] provided similar results for the mAP and PD losses for ranking and their quadratic surrogate. Similarly to our conclusions, the mAP surrogate of [18] is consistent with $\frac{1}{2}r(r+1)$ parameters learned and only low-noise consistent with r parameters learned. Long & Servedio [12] introduced a notion of realizable consistency w.r.t. a function class (they considered linear predictors), which is consistency w.r.t. the function class assuming the data distribution such that labels depend on features deterministically with this dependency being in the correct function class. Ben-David et al. [3] worked in the agnostic setting for binary classification (no assumptions on the underlying \mathcal{D}) and provided guarantees on the error of linear predictors when the margin was bounded by some constant (their work reduces to consistency in the limit case, but is more general).

Conclusion. Differently from the previous approaches, we do not put constraints on the data generating distribution, but instead study the connection between the surrogate and task losses by the means of the calibration function (following [14]), which represents the worst-case scenario. For the quadratic surrogate (11), we can bound the calibration function from below in such a way that the bound is non-trivial in inconsistent settings (differently from [14]). Our bound quantifies the level of inconsistency of a setting (defined by the used surrogate loss, task loss and parametrization of the scores) and allows to analyze when learning with inconsistent surrogates can be beneficial. We illustrate the behavior of our bound for two tasks (multi-class classification and ranking) and show examples of conclusions that our approach can give.

Future work. It would be interesting to combine our quantitative analysis with the constraints on the data distribution, which might give adaptive calibration functions (in analogy to adaptive convergence rates in convex optimization: for example, SAGA [9] has a linear convergence rate for strongly convex objectives and $1/t$ rate for non-strongly convex ones), and with the recent results of Pillaud-Vivien et al. [16] showing that under some low-noise assumptions even slow convergence of the surrogate objective can imply exponentially fast convergence of the task loss.

Acknowledgements

This work was partly supported by Samsung Research, by Samsung Electronics, by the Ministry of Education and Science of the Russian Federation (grant 14.756.31.0001) and by the NSERC Discovery Grant RGPIN-2017-06936.

References

- [1] Ávila Pires, Bernardo, Ghavamzadeh, Mohammad, and Szepesvári, Csaba. Cost-sensitive multiclass classification risk bounds. In *ICML*, 2013.
- [2] Bartlett, Peter L., Jordan, Michael I., and McAuliffe, Jon D. Convexity, classification, and risk bounds. *Journal of the American Statistical Association*, 101(473):138–156, 2006.
- [3] Ben-David, Shai, Loker, David, Srebro, Nathan, and Sridharan, Karthik. Minimizing the misclassification error rate using a surrogate convex loss. 2012.
- [4] Buffoni, David, Gallinari, Patrick, Usunier, Nicolas, and Calauzènes, Clément. Learning scoring functions with order-preserving losses and standardized supervision. In *ICML*, 2011.
- [5] Calauzènes, Clément, Usunier, Nicolas, and Gallinari, Patrick. On the (non-)existence of convex, calibrated surrogate losses for ranking. In *NIPS*, 2012.
- [6] Choromanska, Anna, Agarwal, Alekh, and Langford, John. Extreme multi class classification. In *NIPS Workshop: eXtreme Classification*, 2013.
- [7] Ciliberto, Carlo, Rudi, Alessandro, and Rosasco, Lorenzo. A consistent regularization approach for structured prediction. In *NIPS*, 2016.
- [8] Crammer, Koby and Singer, Yoram. On the algorithmic implementation of multiclass kernel-based vector machines. *Journal of Machine Learning Research (JMLR)*, 2:265–292, 2001.
- [9] Defazio, Aaron, Bach, Francis, and Lacoste-Julien, Simon. SAGA: A fast incremental gradient method with support for non-strongly convex composite objectives. In *NIPS*, 2014.
- [10] Dorn, William S. Duality in quadratic programming. *Quarterly of Applied Mathematics*, 18(2):155–162, 1960.
- [11] Duchi, John C., Mackey, Lester W., and Jordan, Michael I. On the consistency of ranking algorithms. In *ICML*, 2010.
- [12] Long, Phil and Servedio, Rocco. Consistency versus realizable H-consistency for multiclass classification. In *ICML*, 2013.
- [13] Nemirovski, A., Juditsky, A., Lan, G., and Shapiro, A. Robust stochastic approximation approach to stochastic programming. *SIAM Journal on Optimization*, 19(4):1574–1609, 2009.
- [14] Osokin, Anton, Bach, Francis, and Lacoste-Julien, Simon. On structured prediction theory with calibrated convex surrogate losses. In *NIPS*, 2017.
- [15] Pedregosa, Fabian, Bach, Francis, and Gramfort, Alexandre. On the consistency of ordinal regression methods. *Journal of Machine Learning Research (JMLR)*, 18(55):1–35, 2017.
- [16] Pillaud-Vivien, Loucas, Rudi, Alessandro, and Bach, Francis. Exponential convergence of testing error for stochastic gradient methods. In *COLT*, 2018.
- [17] Ramaswamy, Harish G. and Agarwal, Shivani. Convex calibration dimension for multiclass loss matrices. *Journal of Machine Learning Research (JMLR)*, 17(14):1–45, 2016.
- [18] Ramaswamy, Harish G., Agarwal, Shivani, and Tewari, Ambuj. Convex calibrated surrogates for low-rank loss matrices with applications to subset ranking losses. In *NIPS*, 2013.
- [19] Steinwart, Ingo. How to compare different loss functions and their risks. *Constructive Approximation*, 26(2):225–287, 2007.
- [20] Taskar, Ben, Guestrin, Carlos, and Koller, Daphne. Max-margin markov networks. In *NIPS*, 2003.
- [21] Tewari, Ambuj and Bartlett, Peter L. On the consistency of multiclass classification methods. *Journal of Machine Learning Research (JMLR)*, 8:1007–1025, 2007.
- [22] Tsochantaridis, I., Joachims, T., Hofmann, T., and Altun, Y. Large margin methods for structured and interdependent output variables. *Journal of Machine Learning Research (JMLR)*, 6:1453–1484, 2005.
- [23] Zhang, Tong. Statistical analysis of some multi-category large margin classification methods. *Journal of Machine Learning Research (JMLR)*, 5:1225–1251, 2004.
- [24] Zhang, Tong. Statistical behavior and consistency of classification methods based on convex risk minimization. *Annals of Statistics*, 32(1):56–134, 2004.

Supplementary Material (Appendix)

Quantifying Learning Guarantees for Convex but Inconsistent Surrogates

Outline

Section A: Proofs of the two technical lemmas used in Theorem 3.
 Section B: Proof of Theorem 3, which is the main result of this paper.
 Section C: Lower bound on the calibration function for the tree-structured loss.
 Section D: Derivations for the mean average precision loss.

A Technical Lemmas

In this section, we prove two technical lemmas that are used in the proofs of the main theoretical claims of the paper. These two lemmas are the generalizations of the two corresponding lemmas of [14].

Lemma 7 computes the excess of the weighted surrogate risk $\delta\phi$ for the quadratic loss Φ_{quad} (11), which is central to our analysis presented in Section 4. Lemma 7 generalizes Lemma 9 of [14] by removing the assumption of $\text{span}(L) \subseteq \mathcal{F}$. Analogously to Lemma 9 [14], the key property of this result is that the excess $\delta\phi$ is jointly convex w.r.t. the parameters θ and conditional distribution q , which allows further analysis.

Lemma 8 allows to cope with the combinatorial aspect of the calibration function computation. In particular, when the excess of the weighted surrogate risk is convex, Lemma 8 reduces the computation of the calibration function to a set of convex optimization problems, which often can be solved analytically. Note that our Lemma 8 is slightly different from Lemma 10 of Osokin et al. [14] to deal with the difference of the definition of the excess population risk (6).

Lemma 7. *Consider the quadratic surrogate Φ_{quad} (11) defined for a task loss L . Let a subspace of scores $\mathcal{F} \subseteq \mathbb{R}^k$ be parametrized by $\theta \in \mathbb{R}^r$, i.e., $\mathbf{f} = F\theta \in \mathcal{F}$ with $F \in \mathbb{R}^{k \times r}$. Then, the excess of the weighted surrogate loss can be expressed as*

$$\delta\phi_{\text{quad}}(F\theta, q) := \phi_{\text{quad}}(F\theta, q) - \inf_{\theta' \in \mathbb{R}^r} \phi_{\text{quad}}(F\theta', q) = \frac{1}{2k} \|F\theta + P_{\mathcal{F}} Lq\|_2^2,$$

where $P_{\mathcal{F}} := F(F^T F)^{\dagger} F^T$ is the orthogonal projection on the subspace $\mathcal{F} = \text{span}(F)$.

Proof. The proof is almost identical to the proof of Lemma 9 of [14] generalizing it only in the last equality. By the definition of the quadratic surrogate Φ_{quad} (11), we have

$$\begin{aligned} \phi(\mathbf{f}(\theta), q) &= \frac{1}{2k} (\theta^T F^T F \theta + 2\theta^T F^T Lq) + r(q), \\ \theta^* &:= \arg\min_{\theta} \phi(\mathbf{f}(\theta), q) = -(F^T F)^{\dagger} F^T Lq, \\ \delta\phi(\mathbf{f}(\theta), q) &= \frac{1}{2k} (\theta^T F^T F \theta + 2\theta^T F^T Lq + q^T L^T F (F^T F)^{\dagger} F^T Lq) \\ &= \frac{1}{2k} \|F\theta + P_{\mathcal{F}} Lq\|_2^2, \end{aligned}$$

where $r(q)$ denotes the quantity independent of the parameters θ . Note that if the assumption $\text{span}(L) \subseteq \text{span}(F)$ holds we have $P_{\mathcal{F}} L = L$, which is the statement of Lemma 9 [14]. \square

Lemma 8. *For any task loss L , a surrogate loss Φ that is continuous and bounded from below, and a set of scores \mathcal{F} , the calibration function can be lower bounded as*

$$H_{\Phi, L, \mathcal{F}}(\varepsilon) \geq \min_{i \neq j} H_{ij}(\varepsilon), \quad (16)$$

where H_{ij} is defined via minimization of the same objective as (7), but w.r.t. a smaller domain:

$$\begin{aligned}
H_{ij}(\varepsilon) &= \inf_{\mathbf{f}, \mathbf{q}} \delta\phi(\mathbf{f}, \mathbf{q}), \\
&\text{s.t. } \ell_i(\mathbf{q}) \leq \ell_j(\mathbf{q}) - \varepsilon, \\
&\quad \ell_i(\mathbf{q}) \leq \ell_c(\mathbf{q}), \quad \forall c \in \hat{\mathcal{Y}}, \\
&\quad f_j \geq f_c, \quad \forall c \in \hat{\mathcal{Y}}, \\
&\quad \mathbf{f} \in \mathcal{F}, \\
&\quad \mathbf{q} \in \Delta_m.
\end{aligned} \tag{17}$$

Here $\ell_c(\mathbf{q}) := (L\mathbf{q})_c$ is the expected loss if predicting label c . The index i represents a label with the smallest expected loss while the index j represents a label with the largest score.

Proof. We use the notation \mathcal{F}_j to define the set of score vectors \mathbf{f} where the predictor $\text{pred}(\mathbf{f})$ takes the value j , i.e., $\mathcal{F}_j := \{\mathbf{f} \in \mathcal{F} \mid \text{pred}(\mathbf{f}) = j\}$. The union of the sets \mathcal{F}_j , $j \in \hat{\mathcal{Y}}$, equals the whole set \mathcal{F} . Sets \mathcal{F}_j might not contain their boundaries because of the usage of a particular tie-breaking strategy, thus we consider the sets $\overline{\mathcal{F}}_j := \{\mathbf{f} \in \mathcal{F} \mid f_j \geq f_c, \forall c \in \hat{\mathcal{Y}}\}$, which are the closures of \mathcal{F}_j if \mathcal{F}_j are not empty. It also might happen that because of a particular tie-breaking strategy a set \mathcal{F}_j is empty, while the corresponding $\overline{\mathcal{F}}_j$ is not.

If $\mathbf{f} \in \mathcal{F}_j$, i.e. $j = \text{pred}(\mathbf{f})$, then the feasible set of probability vectors \mathbf{q} for which a label i is one of the best possible predictions (i.e. $\delta\ell(\mathbf{f}, \mathbf{q}) = \ell_j(\mathbf{q}) - \ell_i(\mathbf{q}) \geq \varepsilon$) equals

$$\Delta_{m,i,j,\varepsilon} := \{\mathbf{q} \in \Delta_m \mid \ell_i(\mathbf{q}) \leq \ell_c(\mathbf{q}), \forall c \in \hat{\mathcal{Y}}; \ell_j(\mathbf{q}) - \ell_i(\mathbf{q}) \geq \varepsilon\},$$

because $\inf_{\mathbf{f}' \in \mathbb{R}^k} \ell(\mathbf{f}', \mathbf{q}) = \min_{c \in \hat{\mathcal{Y}}} \ell_c(\mathbf{q})$.

The union of the sets $\{\mathcal{F}_j \times \Delta_{m,i,j,\varepsilon}\}_{i,j \in \hat{\mathcal{Y}}, i \neq j}$ exactly equals the feasibility set of the optimization problem (7)-(8) (note that this is not true for the union of the sets $\{\overline{\mathcal{F}}_j \times \Delta_{m,i,j,\varepsilon}\}_{i,j \in \hat{\mathcal{Y}}, i \neq j}$, which can be strictly larger), thus we can rewrite the definition of the calibration function as follows:

$$H_{\Phi, L, \mathcal{F}}(\varepsilon) = \min_{\substack{i,j \in \hat{\mathcal{Y}} \\ i \neq j}} \inf_{\substack{\mathbf{f} \in \mathcal{F}_j, \\ \mathbf{q} \in \Delta_{m,i,j,\varepsilon}}} \delta\phi(\mathbf{f}, \mathbf{q}) \geq \min_{\substack{i,j \in \hat{\mathcal{Y}} \\ i \neq j}} \inf_{\substack{\mathbf{f} \in \overline{\mathcal{F}}_j, \\ \mathbf{q} \in \Delta_{m,i,j,\varepsilon}}} \delta\phi(\mathbf{f}, \mathbf{q}) = \min_{i \neq j} H_{ij}(\varepsilon), \tag{18}$$

which finishes the proof. Note that the inequality of (18) can be not tight only if some \mathcal{F}_j is empty, but the corresponding $\overline{\mathcal{F}}_j$ is not (due to continuity of the function $\delta\phi(\mathbf{f}, \mathbf{q})$, which follows from Lemma 27 of [23]). \square

B Proof of Theorem 3

Theorem 3 (Lower bound on the calibration function $H_{\Phi_{\text{quad}}, L, \mathcal{F}}(\varepsilon)$). *For any task loss L , its quadratic surrogate Φ_{quad} , and a score subspace \mathcal{F} , the calibration function is bounded from below:*

$$H_{\Phi_{\text{quad}}, L, \mathcal{F}}(\varepsilon) \geq \min_{i \neq j} \max_{v \geq 0} \frac{(\varepsilon v - \xi_{ij}(v))_+^2}{2k \|P_{\mathcal{F}} \Delta_{ij}\|_2^2}, \quad \text{where } \xi_{ij}(v) := \|L^\top(v\mathbf{I}_k - P_{\mathcal{F}})\Delta_{ij}\|_\infty, \tag{19}$$

$P_{\mathcal{F}}$ is the orthogonal projection on the subspace \mathcal{F} , $(x)_+^2 := [x > 0]x^2$ is the truncation of the parabola to its right branch and $\Delta_{ij} := \mathbf{e}_i - \mathbf{e}_j \in \mathbb{R}^k$ with $\mathbf{e}_c \in \mathbb{R}^k$ being the c -th column of the identity matrix \mathbf{I}_k . By convention, if both numerator and denominator of (19) equal zero the whole bound equals zero. If only the denominator equals zero then the whole bound equals infinity (the particular pair of i and j is effectively not considered).

Proof. First, let us assume that the score subspace \mathcal{F} is defined as the column space of a matrix $F \in \mathbb{R}^{k \times r}$, i.e., $\mathbf{f}(\boldsymbol{\theta}) = F\boldsymbol{\theta}$. For technical convenience, we can also assume that F is of the full rank, $\text{rank}(F) = r$. Lemma 7 gives us the expression $\delta\phi_{\text{quad}}(F\boldsymbol{\theta}, \mathbf{q}) = \frac{1}{2k} \|F\boldsymbol{\theta} + P_{\mathcal{F}}L\mathbf{q}\|_2^2$ for the excess surrogate, which is jointly convex w.r.t. a conditional probability vector \mathbf{q} and parameters $\boldsymbol{\theta}$.

The optimization problem (7)-(8) is non-convex because the constraint (8) on the excess risk depends of the predictor function $\text{pred}(\mathbf{f})$, see Eq. (1), containing the argmax operation. However, if we

constrain the predictor to output label j , i.e., $f_j \geq f_c, \forall c$, and the label delivering the smallest possible expected loss to be i , i.e., $(L\mathbf{q})_i \leq (L\mathbf{q})_c, \forall c$, the problem becomes convex because all the constraints are linear and the objective is convex. Lemma 8 in Appendix A allows to bound the calibration function with the minimal w.r.t. selected labels i and j optimal value of one of the convex problems, i.e., $H_{\Phi_{\text{quad}}, L, \mathcal{F}}(\varepsilon) \geq \min_{i \neq j} H_{ij}(\varepsilon)$, where $H_{ij}(\varepsilon)$ is defined as follows:

$$\begin{aligned} H_{ij}(\varepsilon) = \min_{\boldsymbol{\theta}, \mathbf{q}} \quad & \frac{1}{2k} \|F\boldsymbol{\theta} + P_{\mathcal{F}}L\mathbf{q}\|_2^2, \\ \text{s.t.} \quad & (L\mathbf{q})_i \leq (L\mathbf{q})_j - \varepsilon, \\ & (L\mathbf{q})_i \leq (L\mathbf{q})_c, \quad \forall c \in \hat{\mathcal{Y}}, \\ & (F\boldsymbol{\theta})_j \geq (F\boldsymbol{\theta})_c, \quad \forall c \in \hat{\mathcal{Y}}, \\ & \mathbf{q} \in \Delta_m. \end{aligned} \quad (20)$$

To obtain a lower bound, we relax (20) by removing some of the constraints and arrive at

$$kH_{ij}(\varepsilon) \geq \min_{\boldsymbol{\theta}, \mathbf{q}} \frac{1}{2} \|F\boldsymbol{\theta} + P_{\mathcal{F}}L\mathbf{q}\|_2^2, \quad (21)$$

$$\text{s.t. } \Delta_{ij}^T L\mathbf{q} \leq -\varepsilon, \quad (22)$$

$$\Delta_{ij}^T F\boldsymbol{\theta} \leq 0, \quad (23)$$

$$\mathbf{1}_m^T \mathbf{q} = 1, \quad (24)$$

$$q_c \geq 0, \quad c = 1, \dots, m. \quad (25)$$

where $\Delta_{ij}^T L\mathbf{q} = (L\mathbf{q})_i - (L\mathbf{q})_j$, $\Delta_{ij}^T F\boldsymbol{\theta} = (F\boldsymbol{\theta})_i - (F\boldsymbol{\theta})_j$, and $\Delta_{ij} = \mathbf{e}_i - \mathbf{e}_j \in \mathbb{R}^k$ with $\mathbf{e}_c \in \mathbb{R}^k$ being a vector with 1 at position c and zeros elsewhere. Note that the relaxation defined by the problem (21)-(25) is tighter than the one used in the proof of Theorem 7 [14, Eq. (25)-(27)], because the latter omitted the simplex constraints (24)-(25).

We now explicitly build a dual problem to the QP (21)-(25). If we used the standard Lagrangian approach we would have to invert the matrix defining the objective, which is difficult. Instead we use the dual formulation of Dorn [10, Page 160], which allows to build a dual without inverting any matrices.⁹ For the problem (21)-(25), this dual can be written as follows:

$$kH_{ij}(\varepsilon) \geq \max_{\boldsymbol{\theta}, \mathbf{q}, v_F \geq 0, v_L \geq 0, u} - \frac{1}{2} \|F\boldsymbol{\theta} + P_{\mathcal{F}}L\mathbf{q}\|_2^2 + v_L \varepsilon + u, \quad (26)$$

$$-v_L L^T \Delta_{ij} + u \mathbf{1}_m - L^T P_{\mathcal{F}}L\mathbf{q} - L^T F\boldsymbol{\theta} \leq \mathbf{0}_m, \quad (27)$$

$$-v_F F^T \Delta_{ij} - F^T L\mathbf{q} - F^T F\boldsymbol{\theta} = \mathbf{0}_r. \quad (28)$$

From the equality (28), we can express $F^T L\mathbf{q} = -v_F F^T \Delta_{ij} - F^T F\boldsymbol{\theta}$ and substitute it in the objective (26) and inequality (27). Using the identities $P_{\mathcal{F}} = F(F^T F)^{-1} F^T$ and $P_{\mathcal{F}} F = F$, we can exclude variables $\boldsymbol{\theta}, \mathbf{q}$ and get a simpler bound. Note that this step leads to a valid lower bound because for any $v_F \geq 0$ there exist feasible values of variables \mathbf{q} and $\boldsymbol{\theta}$ (we can take simply $\mathbf{q} = \mathbf{0}$, $\boldsymbol{\theta} = -v_F (F^T F)^{-1} F^T \Delta_{ij}$). The new bound depends on the three variables only:

$$kH_{ij}(\varepsilon) \geq \max_{v_F \geq 0, v_L \geq 0, u} - \frac{1}{2} v_F^2 \Delta_{ij}^T P_{\mathcal{F}} \Delta_{ij} + v_L \varepsilon + u, \quad (29)$$

$$-v_L L^T \Delta_{ij} + u \mathbf{1}_m + v_F L^T P_{\mathcal{F}} \Delta_{ij} \leq \mathbf{0}_m. \quad (30)$$

⁹Here we show the dual of Dorn [10] for the exact combination of constraints we are using. In the dual formulation, \mathbf{v} and \mathbf{u} are the extra variables corresponding to the inequality and equality constraints, respectively.

<p style="text-align: center;">The primal problem</p> $\min_{\mathbf{q} \geq 0, \boldsymbol{\theta}} \frac{1}{2} (\mathbf{q}^T \boldsymbol{\theta}^T) \begin{pmatrix} H_{qq} & H_{q\theta} \\ H_{q\theta}^T & H_{\theta\theta} \end{pmatrix} \begin{pmatrix} \mathbf{q} \\ \boldsymbol{\theta} \end{pmatrix},$ <p style="text-align: center;">s.t. $A_q \mathbf{q} + A_\theta \boldsymbol{\theta} \geq \mathbf{b}$</p> <p style="text-align: center;">$C_q \mathbf{q} + C_\theta \boldsymbol{\theta} = \mathbf{d}$</p>	<p style="text-align: center;">The dual problem</p> $\max_{\mathbf{q}, \boldsymbol{\theta}, v \geq 0, \mathbf{u}} - \frac{1}{2} (\mathbf{q}^T \boldsymbol{\theta}^T) \begin{pmatrix} H_{qq} & H_{q\theta} \\ H_{q\theta}^T & H_{\theta\theta} \end{pmatrix} \begin{pmatrix} \mathbf{q} \\ \boldsymbol{\theta} \end{pmatrix} + \mathbf{b}^T \mathbf{v} + \mathbf{d}^T \mathbf{u},$ <p style="text-align: center;">s.t. $A_q^T \mathbf{v} + C_q^T \mathbf{u} - H_{qq} \mathbf{q} - H_{q\theta} \boldsymbol{\theta} \leq \mathbf{0}$</p> <p style="text-align: center;">$A_\theta^T \mathbf{v} + C_\theta^T \mathbf{u} - H_{q\theta}^T \mathbf{q} - H_{\theta\theta} \boldsymbol{\theta} = \mathbf{0}$</p>
--	---




Figure 2: *Left*: An example of the tree-structured loss for the task of multi-class classification. *Right*: Illustration of the proof of Lemma 9 (best viewed in color). The thin gray and brown lines show the absolute values of the components of the vector $L_{\text{tree}}^T(v\mathbf{I} - P_{\mathcal{F}_{\text{tree}}})\Delta_{ij}$ as functions of v . The bold blue and green lines correspond to the components at which the maximum value is achieved. The bold red line shows the resulting norm $\|L_{\text{tree}}^T(v\mathbf{I} - P_{\mathcal{F}_{\text{tree}}})\Delta_{ij}\|_{\infty}$.

First, consider the case $\Delta_{ij}^T P_{\mathcal{F}} \Delta_{ij} = \|P_{\mathcal{F}} \Delta_{ij}\|_2^2 \neq 0$. Given that $v_F \geq 0$ we can change the variables by introducing $\hat{v}_F := v_F \|P_{\mathcal{F}} \Delta_{ij}\|_2^2$, $v := v_L / v_F$, $\hat{u} := u / v_F$ after which we get

$$kH_{ij}(\varepsilon) \geq \frac{1}{\|P_{\mathcal{F}} \Delta_{ij}\|_2^2} \max_{\hat{v}_F \geq 0, \hat{v}_L \geq 0, \hat{u}} -\frac{1}{2} \hat{v}_F^2 + \hat{v}_F(v\varepsilon + \hat{u}), \quad (31)$$

$$-vL^T \Delta_{ij} + \hat{u} \mathbf{1}_m + L^T P_{\mathcal{F}} \Delta_{ij} \leq \mathbf{0}_m. \quad (32)$$

The global minimum of this function w.r.t. the variable \hat{v}_F can be found analytically: if $v\varepsilon + \hat{u} \geq 0$ it equals $\frac{1}{2\|P_{\mathcal{F}} \Delta_{ij}\|_2^2} (v\varepsilon + \hat{u})^2$, and zero otherwise. The constraint (32) on \hat{u} can be substituted with $\hat{u} = -\|L^T(v\mathbf{I}_k - P_{\mathcal{F}})\Delta_{ij}\|_{\infty} =: -\xi_{ij}(v)$, because we always consider both $H_{ij}(\varepsilon)$ and $H_{ji}(\varepsilon)$ when bounding the calibration function.

Now, consider the boundary case of $\Delta_{ij}^T P_{\mathcal{F}} \Delta_{ij} = \|P_{\mathcal{F}} \Delta_{ij}\|_2^2 = 0$. The problem (29)-(30) becomes $\frac{1}{2} \max_{v \geq 0} v(\varepsilon + \min L^T \Delta_{ij})$ implying that the objective equals 0 if $\varepsilon + \min L^T \Delta_{ij} \leq 0$. Otherwise, the objective equals $+\infty$, which corresponds to the in-feasibility of the constraint (22) of the primal problem. Note that because we always consider both $H_{ij}(\varepsilon)$ and $H_{ji}(\varepsilon)$ when bounding the calibration function we can substitute $v \min(L^T \Delta_{ij})$ with $-\xi_{ij}(v)$. \square

C Lower Bound on the Calibration Function for the Tree-Structured Loss

In this section, we compute the lower bound on the calibration function for the tree-structured loss defined in Section 5.1.

Lemma 9. *For a particular consistency depth s_0 and for the corresponding subspace $\mathcal{F}_{\text{tree}, s_0}$, the projection operator $P_{\mathcal{F}_{\text{tree}, s_0}}$ at Δ_{ij} is computed as*

$$P_{\mathcal{F}_{\text{tree}, s_0}} \Delta_{ij} = \begin{cases} 0, & i \in B(j), \\ \frac{1}{|B(j)|} \left(\sum_{k \in B(i)} \mathbf{e}_k - \sum_{k \in B(j)} \mathbf{e}_k \right) & i \notin B(j). \end{cases} \quad (33)$$

The vectors $\xi_{ij}(v)$ are computed as

$$\xi_{ij}(v) = \begin{cases} v\rho(i, j) & i \in B(j), \\ \max\{|(v-1)\rho(i, j)|, |v(\rho(i, j) - \eta) - (\rho(i, j) - \bar{\rho})|\} & i \notin B(j), \end{cases} \quad (34)$$

for $\eta := \max_{c \in B(i)} \rho(i, c)$ and $\bar{\rho} := \frac{1}{|B(j)|} \sum_{c \in B(j)} \rho(i, c)$. Finally, the following lower bound of the calibration function for the loss L_{tree} , its quadratic surrogate Φ_{quad} and the score subspace $\mathcal{F}_{\text{tree}, s_0}$ holds:

$$H_{\Phi_{\text{quad}}, L_{\text{tree}}, \mathcal{F}_{\text{tree}, s_0}}(\varepsilon) \geq [\varepsilon > \eta] \frac{(\nu - \bar{\rho})^2}{(\nu - \frac{\eta}{2})^2} \frac{(\varepsilon - \frac{\eta}{2})_+^2}{4b}, \quad (35)$$

where $\nu := \min_{c \notin B(i)} \rho(i, c) > \eta$ and b is the number of blocks when the tree is cut at the depth s_0 .

Proof. For brevity, we shortcut the notation F_{tree, s_0} to F , $\mathcal{F}_{\text{tree}, s_0}$ to \mathcal{F} and L_{tree} to L . First, we compute the projection operator $P_{\mathcal{F}}\mathbf{e}_i$ and the lower-bound denominator $2k\|P_{\mathcal{F}}\Delta_{ij}\|_2^2$. Recall, that the subspace of allowed scores \mathcal{F} defined as $\text{span}\{\sum_{l \in B(j)} \mathbf{e}_l | j = 1, \dots, k\}$ is of dimension b . The vector \mathbf{e}_i is orthogonal to the $b - 1$ different vectors $\sum_{l \in B(j)} \mathbf{e}_l$, $j \neq B(i)$, thus the projection $P_{\mathcal{F}}\mathbf{e}_i$ equals the projection of \mathbf{e}_i on the vector $\sum_{l \in B(i)} \mathbf{e}_l$, which equals $\frac{1}{s} \sum_{l \in B(i)} \mathbf{e}_l$ with $s := B(i) = \frac{k}{b}$. The projection square norm $\|P_{\mathcal{F}}\Delta_{ij}\|_2^2$ equals $\frac{2}{s} = \frac{2b}{k}$ if $i \notin B(j)$ and 0 if $i \in B(j)$.

Next, we compute $\xi_{ij}(v)$ defined as $\|L^\top(v\mathbf{I} - P_{\mathcal{F}})\Delta_{ij}\|_\infty$. By definition of the loss function, the element of the loss matrix $L_{ci} = (L\mathbf{e}_i)_c$ equals the tree distance from the leaf i to the leaf c . The projection operator $P_{\mathcal{F}}\mathbf{e}_i$ equals the vector $\frac{1}{s} \sum_{l \in B(i)} \mathbf{e}_l$, therefore $(LP_{\mathcal{F}}\mathbf{e}_i)_c$ is equal to the average tree distance from the elements of the block $B(i)$ to c : $\frac{1}{s} \sum_{l \in B(i)} L_{lc}$, which we denote by $\bar{\rho}(i, c)$. Note that the average distance $\bar{\rho}(i, c)$ is equal for all the leaves c that belong to the same block $B(c)$. With this notation, we have the following equality:

$$\|L^\top(v\mathbf{I} - P_{\mathcal{F}})\Delta_{ij}\|_\infty = \max_{c \in \mathcal{Y}} |v(\rho(i, c) - \rho(j, c)) - (\bar{\rho}(i, c) - \bar{\rho}(j, c))|. \quad (36)$$

On the right-hand side, each component is the absolute value of a linear in v function, which equals zero at $v = \frac{\bar{\rho}(i, c) - \bar{\rho}(j, c)}{\rho(i, c) - \rho(j, c)}$ and the absolute value of the slope equals $|\rho(i, c) - \rho(j, c)|$. We consider the cases when the labels i and j are in the same and different blocks separately.

If i and j are in the same block we have that $\bar{\rho}(i, c) = \bar{\rho}(j, c)$ and, by the reverse triangle inequality, $|\rho(i, c) - \rho(j, c)| \leq \rho(i, j)$ with the equality holding for $c = i$ or $c = j$, which implies that $\xi_{ij}(v) = \|(L^\top(v\mathbf{I} - P_{\mathcal{F}})\Delta_{ij})\|_\infty = v\rho(i, j)$ for $i \in B(j)$.

Now, we study the second case where i and j are in different blocks. We first show that

$$|\bar{\rho}(i, c_1) - \bar{\rho}(j, c_1)| \leq |\bar{\rho}(i, c_2) - \bar{\rho}(j, c_2)| \quad \text{if } c_1 \notin B(i) \cup B(j) \text{ and } c_2 \in B(i) \cup B(j). \quad (37)$$

This inequality is crucial for the proof and holds due to the restriction on tree weights and node degrees.

In this paragraph, we will show that the left-hand side of the inequality (37) achieves its maximum when $c_1 \notin B(i) \cup B(j)$ is in the block closest to $B(j)$ (or, due to the loss symmetries, in the block closest to $B(i)$). If the lowest common ancestor of i and j is not an ancestor of c_1 the difference of the average distances equals zero due to equality of the paths from the lowest common ancestor to i and j . Otherwise, there exists $c_1 \notin B(i) \cup B(j)$ such that the lowest common ancestor of i and j is an ancestor of c_1 . Then, $\bar{\rho}(j, c_1)$ is minimized and $\bar{\rho}(i, c_1)$ is simultaneously maximized for a component c_1 closest to the block $B(j)$. In this case, the left-hand side maximum value equals $\bar{\rho}(i, j) - \min_{c \notin B(j)} \bar{\rho}(j, c)$ because $\bar{\rho}(i, j) = \bar{\rho}(i, c_1)$.

The right-hand side of the inequality (37) is the same for any choice of $c_2 \in B(i) \cup B(j)$ and is equal to $\bar{\rho}(i, j) - \bar{\rho}(j, c_2)$ for some $c_2 \in B(j)$. Since the average distance within the block is smaller than the average distance to any node outside of the block, i.e., $\min_{c \notin B(j)} \bar{\rho}(j, c) \geq \bar{\rho}(j, c_2)$ for $c_2 \in B(j)$, the inequality (37) holds. The same arguments also show that

$$|\rho(i, c_1) - \rho(j, c_1)| \leq |\rho(i, c_2) - \rho(j, c_2)| \quad \text{if } c_1 \notin B(i) \cup B(j) \text{ and } c_2 \in B(i) \cup B(j). \quad (38)$$

Recall that in our case the infinity norm in (36) equals the component-wise maximum of the absolute values of the linear functions of v . We will show below that for a small enough v the maximum is achieved at the components that have the smallest slope $|\rho(i, c) - \rho(j, c)|$ among the ones with the largest offset $|\bar{\rho}(i, c) - \bar{\rho}(j, c)|$ and from some point for larger values of v the maximum is achieved at the components with the steepest slope (see Figure 2 right for the illustration).

Consider a leaf $c_2 \in B(i)$ farthest from the leaf i , i.e., $c_2 \in \arg\max_{c \in B(i)} \rho(i, c)$ (defines the green line in Figure 2 right). The offset $|\bar{\rho}(i, c_2) - \bar{\rho}(j, c_2)|$ is the same for all $c_2 \in B(i)$ and, by (37), is larger than the offsets of the components $c_1 \notin B(i) \cup B(j)$. The slope $|\rho(i, c_2) - \rho(j, c_2)|$ is the smallest among the components in $B(i) \cup B(j)$. The component c_2 of $L^\top(v\mathbf{I} - P_{\mathcal{F}})\Delta_{ij}$ equals zero for $v_{c_2}^* := \frac{\bar{\rho}(j, c_2) - \bar{\rho}(i, c_2)}{\rho(j, c_2) - \rho(i, c_2)} = \frac{\rho(i, j) - \bar{\rho}(i, c_2)}{\rho(i, j) - \rho(i, c_2)}$, where $v_{c_2}^* > 1$ by definition of c_2 , i.e., because $\rho(i, c_2)$ is the maximal distance, which is not smaller than the average distance $\bar{\rho}(i, c_2)$. Finally, for $v \leq 1$ this component has higher values than the values of the components $c \notin B(i) \cup B(j)$. Indeed, the latter

are equal to zero at $v = 1$ and have smaller offset at $v = 0$ (thin brown lines in Figure 2 right for $v < 1$).

The component i of $L^\top(v\mathbf{I} - P_{\mathcal{F}})\Delta_{ij}$ has the steepest slope $|\rho(i, j) - \rho(i, i)| = \rho(i, j)$ and the same offset as in the previous paragraph $|\bar{\rho}(i, i) - \bar{\rho}(i, j)| = |\bar{\rho}(i, c_2) - \bar{\rho}(j, c_2)|$ (defines the blue line in Figure 2 right). The component equals zero for $v_i^* := \frac{\bar{\rho}(j, i) - \bar{\rho}(i, i)}{\rho(j, i) - \rho(i, i)} = \frac{\rho(j, i) - \bar{\rho}(i, i)}{\rho(j, i)}$, where $v_i^* \leq 1$. As a result, the component i has higher values than the components $c \notin B(i) \cup B(j)$, since they have smaller slope $|\rho(i, c) - \rho(i, c)|$ (due to the inequality (38)) and equal zero for $v = 1$ (thin brown lines in Figure 2 right for $v > 1$).

Since all the components $c \in B(i) \cup B(j)$ have the same offset, the maximum is achieved either at c_2 or at i :

$$\|L^\top(v\mathbf{I} - P_{\mathcal{F}})\Delta_{ij}\|_\infty = \max\{|v\rho(i, j) - (\rho(i, j) - \bar{\rho})|, |v(\rho(i, j) - \eta) - (\rho(i, j) - \bar{\rho})|\}, \quad (39)$$

where $\eta := \rho(i, c_2)$ is the maximal distance within a block and $\bar{\rho} := \bar{\rho}(i, c_2)$ is the average distance within a block.

Next, we compute $\max_{v \geq 0}(\varepsilon v - \xi_{ij}(v))_+^2$. If i and j are in the same block we have $(\varepsilon v - \xi_{ij}(v))_+^2 = (v(\varepsilon - \rho(i, j)))_+^2$, which equals zero for $\varepsilon \leq \rho(i, j)$ and $+\infty$ otherwise. If i and j are in different blocks we have the maximum value of $(\varepsilon v - \xi_{ij}(v))_+^2$ equal to $+\infty$ when $\varepsilon > \rho(i, j)$. In the case when $\varepsilon \leq \rho(i, j)$, the maximum is achieved at the intersection point $v\rho(i, j) - (\rho(i, j) - \bar{\rho}) = -v(\rho(i, j) - \eta) + (\rho(i, j) - \bar{\rho})$, $v = \frac{2(\rho(i, j) - \bar{\rho})}{2\rho(i, j) - \eta}$. The maximum value is positive if and only if $\varepsilon > \frac{\eta}{2}$, so for $\varepsilon \leq \rho(i, j)$ we obtain

$$\max_{v \geq 0}(\varepsilon v - \xi_{ij}(v))_+^2 = \frac{(\rho(i, j) - \bar{\rho})^2}{(\rho(i, j) - \frac{\eta}{2})^2}(\varepsilon - \frac{\eta}{2})_+^2 \quad (40)$$

and $+\infty$ otherwise.

Finally, to get the actual lower bound on the calibration function, we compute the minimum with respect to all labels $\min_{i \neq j} \max_{v \geq 0}(\varepsilon v - \xi_{ij}(v))_+^2$. When i and j are in the same block, they deliver minimum value 0 for $\varepsilon \leq \rho(i, j)$ and the maximum value of $\rho(i, j)$ within a block equals η by definition of η . For $\varepsilon > \eta$, the minimum is delivered by i and j in different blocks. For the average distance within the block, we have $\bar{\rho} \geq \frac{\eta}{2}$ for the trees with the number of children and the weights of edges equal at the same depth level, therefore the outer minimum w.r.t. i and j is achieved at the smallest distance between two blocks $\nu := \min_{i \notin B(j)} \rho(i, j) > \eta$. As a result, we obtain the bound

$$H_{\Phi_{\text{quad}}, L_{\text{tree}}, \mathcal{F}_{\text{tree}, s_0}}(\varepsilon) \geq [\varepsilon > \eta] \frac{(\nu - \bar{\rho})^2}{(\nu - \frac{\eta}{2})^2} \frac{(\varepsilon - \frac{\eta}{2})_+^2}{4b}, \quad (41)$$

which completes the proof. \square

In the next lemma, we compute the quantities $\eta, \bar{\rho}, \nu$ using the tree weights $\{\frac{1}{2}\alpha_s\}_{s=0}^{D-1}$ to finish the computation of the bound (14) of the main paper.

Lemma 10. *For a particular consistency depth s_0 and the corresponding subspace $\mathcal{F}_{\text{tree}, s_0}$, the maximum distance within an arbitrary block η_{s_0} , the minimum distance between a leaf in a block and a leaf outside the block ν_{s_0} and the average distance within a block $\bar{\rho}_{s_0}$ can be computed as follows:*

$$\eta_{s_0} = \max_{i \in B(j)} \rho(i, j) = \sum_{s=s_0}^{D-1} \alpha_s \quad (42)$$

$$\nu_{s_0} = \min_{i \notin B(j)} \rho(i, j) = \sum_{s=s_0-1}^{D-1} \alpha_s \quad (43)$$

$$\bar{\rho}_{s_0} = \frac{1}{|B(j)|} \sum_{i \in B(j)} \rho(i, j) = \sum_{s=s_0}^{D-1} \alpha_s \frac{(\prod_{s'=s_0}^s d_{s'}) - 1}{\prod_{s'=s_0}^s d_{s'}}. \quad (44)$$

Proof. The expressions for η_{s_0} and ν_{s_0} immediately follow from the definition of the distance $\rho(i, j)$.

To obtain the expression for $\bar{\rho}_{s_0}$, we rewrite the distance $\rho(i, j)$ between leaves i and j in the same block $B(j)$ as the weighted sum of indicators:

$$\rho(i, j) = \sum_{s=s_0}^{D-1} \alpha_s [\text{path from } i \text{ to } j \text{ contains an edge of depth } s]. \quad (45)$$

Then, we fix a leaf j and compute the number of paths from j to the leaves in $B(j)$ that contain an edge of depth s . Such paths go through the same node (the ancestor of j at the depth s) on the way up from node j and go through one of $\prod_{s'=s_0}^s d_{s'} - 1$ possible nodes at the depth s on the way down. From each node of depth s , the path can further go to one of $\prod_{s'=s+1}^{D-1} d_{s'}$ leaves on the way down. Therefore, there are $(\prod_{s'=s_0}^s d_{s'} - 1) \left(\prod_{s'=s+1}^{D-1} d_{s'} \right)$ paths that contain an edge of depth s .

Next, we rewrite $\sum_{i \in B(j)} \rho(i, j)$ using the indicator notation and compute the sum:

$$\sum_{i \in B(j)} \rho(i, j) = \sum_{s=s_0}^{D-1} \sum_{i \in B(j)} \alpha_s [\text{path from } j \text{ to } i \text{ contains an edge of depth } s] \quad (46)$$

$$= \sum_{s=s_0}^{D-1} \alpha_s \left(\prod_{s'=s_0}^s d_{s'} - 1 \right) \left(\prod_{s'=s+1}^{D-1} d_{s'} \right). \quad (47)$$

Since the number of leaves in a block is $\prod_{s'=s_0}^{D-1} d_{s'}$, we have

$$\bar{\rho}_{s_0} = \frac{1}{|B(j)|} \sum_{i \in B(j)} \rho(i, j) = \sum_{s=s_0}^{D-1} \alpha_s \frac{(\prod_{s'=s_0}^s d_{s'} - 1)}{\prod_{s'=s_0}^s d_{s'}}, \quad (48)$$

which finishes the proof. \square

Note that for tree-depth $D = 2$ the minimum ν_1 equals $\alpha_0 + \alpha_1 = 1$. As a result, our calibration function lower bound coincides with the exact calibration function from [14].

D Derivations for the Mean Average Precision Loss

In this section, we prove several statements about F_{sort} and L_{mAP} , which are used in Section 5.2.

Lemma 11. *The matrix $F_{\text{sort}}^\top F_{\text{sort}}$ has the following form:*

$$(F_{\text{sort}}^\top F_{\text{sort}})_{pq} = \begin{cases} (r-1)!H_{r,2}, & p = q, \\ (r-2)!(H_{r,1}^2 - H_{r,2}), & p \neq q, \end{cases} \quad (49)$$

where $H_{n,m} := \sum_{k=1}^n \frac{1}{k^m}$ is the generalized harmonic number of order m of n . As a result, for distinct permutations π and ω , the square norm of the projection is equal to

$$\|P_{\mathcal{F}_{\text{sort}}} \Delta_{\pi\omega}\|_2^2 = \frac{1}{(r-2)!(rH_{r,2} - H_{r,1}^2)} \sum_{p=1}^r \left(\frac{1}{\pi(p)} - \frac{1}{\omega(p)} \right)^2. \quad (50)$$

The condition number $\kappa(F_{\text{sort}})$ equals $\frac{\sqrt{r-1}H_{r,1}}{\sqrt{rH_{r,2} - H_{r,1}^2}}$.

Proof. By definition, $(F_{\text{sort}}^\top F_{\text{sort}})_{pq} = \sum_{\sigma \in S_r} \frac{1}{\sigma(p)\sigma(q)}$. We can rewrite the sum as the sum over the permutations with fixed values $\sigma(p)$ and $\sigma(q)$ and then sum over the fixed values. Therefore, the sum is equal to $(r-1)!H_{r,2}$ when $p = q$ and is equal to $(r-2)!(H_{r,1}^2 - H_{r,2})$ otherwise.

We now have $F_{\text{sort}}^\top F_{\text{sort}} = (r-2)!(rH_{r,2} - H_{r,1}^2)\mathbf{I}_r + (r-2)!(H_{r,1}^2 - H_{r,2})\mathbf{1}\mathbf{1}^\top$. The Sherman-Woodbury formula for the matrix inversion gives us the sum of the scalar matrix $\frac{1}{(r-2)!(rH_{r,2} - H_{r,1}^2)}\mathbf{I}_r$ and the constant matrix. Since $\mathbf{1}^\top F_{\text{sort}}^\top \Delta_{\pi\omega} = \sum_{p=1}^r \left(\frac{1}{\pi(p)} - \frac{1}{\omega(p)} \right) = 0$, the square norm of the projection equals $\frac{1}{(r-2)!(rH_{r,2} - H_{r,1}^2)} \Delta_{\pi\omega}^\top F_{\text{sort}}^\top F_{\text{sort}} \Delta_{\pi\omega} = \frac{1}{(r-2)!(rH_{r,2} - H_{r,1}^2)} \sum_{p=1}^r \left(\frac{1}{\pi(p)} - \frac{1}{\omega(p)} \right)^2$.

The condition number of F_{sort} equals the square root of the ratio between the maximal and minimal eigenvalues of $F_{\text{sort}}^T F_{\text{sort}}$. Subtracting $(r-2)!(rH_{r,2} - H_{r,1}^2)\mathbf{I}_r$ from $F_{\text{sort}}^T F_{\text{sort}}$ we get a matrix of rank 1, which means that $r-1$ eigenvalues of $F_{\text{sort}}^T F_{\text{sort}}$ equal $(r-2)!(rH_{r,2} - H_{r,1}^2)$. The remaining eigenvalue corresponds to the eigenvector $\mathbf{1}$ and equals $(r-1)!H_{r,1}^2$. With these eigenvalues we get the condition number $\kappa(F_{\text{sort}}) = \frac{\sqrt{r-1}H_{r,1}}{\sqrt{rH_{r,2} - H_{r,1}^2}}$. \square

Lemma 12. *The matrix $L_{\text{mAP}}^T F_{\text{sort}}$ has the following form:*

$$\left(L_{\text{mAP}}^T F_{\text{sort}} \right)_{y,p} = \begin{cases} \alpha(|\mathbf{y}|), & y_p = 1, \\ \beta(|\mathbf{y}|), & y_p = 0. \end{cases} \quad (51)$$

That is, for each ground-truth value \mathbf{y} the matrix row components have only two values that depend on the Hamming norm $|\mathbf{y}| := \sum_{p=1}^r y_p$, specifically:

$$\alpha(|\mathbf{y}|) = \mathfrak{A}_r \left(1 - \frac{|\mathbf{y}|-1}{r-2} \left(1 - \frac{r}{|\mathbf{y}|(r-1)} \right) \right) - \mathfrak{B}_r \left(\frac{3}{2} \frac{|\mathbf{y}|-1}{r-2} - \frac{r-|\mathbf{y}|}{|\mathbf{y}|} \right) - \mathfrak{C}_r \left(\frac{r-|\mathbf{y}|}{|\mathbf{y}|(r-1)} \right), \quad (52)$$

$$\beta(|\mathbf{y}|) = \mathfrak{A}_r \left(1 - \frac{|\mathbf{y}|-1}{r-2} \right) - \mathfrak{B}_r \left(1 - \frac{3}{2} \frac{|\mathbf{y}|-1}{r-2} \right). \quad (53)$$

Here $\mathfrak{A}_r = (r-1)!H_{r,1}$, $\mathfrak{B}_r = (r-2)!(H_{r,1}^2 - H_{r,2})$, $\mathfrak{C}_r = (r-1)!H_{r,2}$. As a result, for permutations π and ω , we obtain

$$\left(L_{\text{mAP}}^T P_{\mathcal{F}_{\text{sort}}} \Delta_{\pi\omega} \right)_{\mathbf{y}} = \gamma(|\mathbf{y}|) \left((F_{\text{sort}}\mathbf{y})_{\pi} - (F_{\text{sort}}\mathbf{y})_{\omega} \right), \quad (54)$$

where $\gamma(p) = \frac{\alpha(p) - \beta(p)}{(r-2)!(rH_{r,2} - H_{r,1}^2)}$.

Proof. For brevity, here we denote F_{sort} by F , $\mathcal{F}_{\text{sort}}$ by \mathcal{F} and L_{mAP} by L . Following the definitions of L and F , we explicitly compute the components of $L^T F$:

$$\left(L^T F \right)_{\mathbf{y},s} = \sum_{\sigma \in S_r} \left(1 - \frac{1}{|\mathbf{y}|} \sum_{p=1}^r \sum_{q=1}^p \frac{y_{\sigma^{-1}(p)} y_{\sigma^{-1}(q)}}{p} \right) \frac{1}{\sigma(s)}. \quad (55)$$

There are exactly $(r-1)!$ permutations with one fixed element, so we have $\sum_{\sigma \in S_r} \frac{1}{\sigma(s)} = (r-1)! \sum_{p=1}^r \frac{1}{p} = (r-1)!H_{r,1} =: \mathfrak{A}_r$. To compute the remaining part of (55), we group the permutation values $\sigma(k) = t$ by each $t = 1, \dots, r$ and move the sum over permutations inside the bracket:

$$-\frac{1}{|\mathbf{y}|} \sum_{\sigma \in S_r} \sum_{p=1}^r \sum_{q=1}^p \frac{y_{\sigma^{-1}(p)} y_{\sigma^{-1}(q)}}{p \sigma(s)} = -\frac{1}{|\mathbf{y}|} \sum_{t=1}^r \sum_{p=1}^r \sum_{q=1}^p \sum_{\sigma \in S_r, \sigma(s)=t} \frac{y_{\sigma^{-1}(p)} y_{\sigma^{-1}(q)}}{pt}. \quad (56)$$

Next, we compute the inner sum $\sum_{\sigma \in S_r, \sigma(s)=t} y_{\sigma^{-1}(p)} y_{\sigma^{-1}(q)}$. We rewrite the sum as the sum over inverse permutations:

$$\sum_{\sigma \in S_r, \sigma(s)=t} y_{\sigma^{-1}(p)} y_{\sigma^{-1}(q)} = \sum_{\pi \in S_r, \pi(t)=s} y_{\pi(p)} y_{\pi(q)}. \quad (57)$$

The number of positive terms is different for the two cases of $y_s = 0$ and $y_s = 1$. For $y_s = 0$, using the Iverson brackets the sum can be rewritten as follows:

$$\sum_{\pi \in S_r, \pi(t)=s} y_{\pi(p)} y_{\pi(q)} = [p \neq t] ([q < p \ \& \ q \neq t] |\mathbf{y}|(|\mathbf{y}|-1)(r-3)! + [q = p] |\mathbf{y}|(r-2)!). \quad (58)$$

We then sum the expression over $q = 1, \dots, p$:

$$\sum_{q=1}^p [p \neq t] ([q < p, q \neq t] |\mathbf{y}|(|\mathbf{y}|-1)(r-3)! + [q = p] |\mathbf{y}|(r-2)!) = \quad (59)$$

$$[p \neq t] ((p-1 - [t < p]) |\mathbf{y}|(|\mathbf{y}|-1)(r-3)! + |\mathbf{y}|(r-2)!). \quad (60)$$

Finally, we multiply the expression by $\frac{-1}{|\mathbf{y}|^{pt}}$ and compute the sum over p and t :

$$-\frac{1}{|\mathbf{y}|} \sum_{t=1}^r \sum_{p=1}^r \frac{1}{t^p} [p \neq t] ((p-1 - [t < p]) |\mathbf{y}| (|\mathbf{y}| - 1)(r-3)! + |\mathbf{y}|(r-2)!) = \quad (61)$$

$$-\mathfrak{A}_r \frac{|\mathbf{y}| - 1}{r-2} + \mathfrak{B}_r \left(1 - \frac{3}{2} \frac{|\mathbf{y}| - 1}{r-2}\right). \quad (62)$$

Combining with $\sum_{\sigma \in S_r} \frac{1}{\sigma(s)} = \mathfrak{A}_r$ we obtain the desired expression for $\beta(|\mathbf{y}|)$.

Now, consider the case of $y_s = 1$. Again, we rewrite the sum as

$$\sum_{\pi \in S_r, \pi(t)=s} y_{\pi(p)} y_{\pi(q)} = [p = t] ([q < p] (|\mathbf{y}| - 1)(r-2)! + [q = p](r-1)!) \quad (63)$$

$$+ [p \neq t] ([q < p \& q = t] (|\mathbf{y}| - 1)(r-2)!) \quad (64)$$

$$+ [p \neq t] ([q < p \& q \neq t] (|\mathbf{y}| - 1)(|\mathbf{y}| - 2)(r-3)!) \quad (65)$$

$$+ [p \neq t] ([q = p] (|\mathbf{y}| - 1)(r-2)!), \quad (66)$$

sum it over q ,

$$\sum_{q=1}^p \sum_{\pi \in S_r, \pi(t)=s} y_{\pi(p)} y_{\pi(q)} = [p = t] ((p-1)(|\mathbf{y}| - 1)(r-2)! + (r-1)!) \quad (67)$$

$$+ [p \neq t] ([t < p] (|\mathbf{y}| - 1)(r-2)!) \quad (68)$$

$$+ [p \neq t] ((p-1 - [t < q]) (|\mathbf{y}| - 1)(|\mathbf{y}| - 2)(r-3)!) \quad (69)$$

$$+ [p \neq t] (|\mathbf{y}| - 1)(r-2)!, \quad (70)$$

and obtain the desired expression by multiplying the latter by $\frac{-1}{|\mathbf{y}|^{pt}}$ and summing over p and t . The last step of the computation is completely analogous to the case $y_k = 0$ and we omit it for brevity.

To compute $(L^T P_{\mathcal{F}} \Delta_{\pi\omega})_{\mathbf{y}}$, we note that for all permutations π and ω it holds that $\mathbf{1}^T F^T \Delta_{\pi\omega} = 0$.

According to Lemma 11 and the Sherman-Woodbury formula, $(F^T F)^{-1}$ is the sum of the scalar matrix $\frac{\mathbf{I}_r}{(r-2)!(rH_{r,2} - H_{r,1}^2)}$ and the multiple of the rank one matrix $\mathbf{1}\mathbf{1}^T$. After the multiplication

on $F^T \Delta_{\pi\omega}$, the second term vanishes, so we get $(F^T F)^{-1} F^T \Delta_{\pi\omega} = \frac{\sum_{p=1}^r \frac{1}{\pi(p)} - \frac{1}{\omega(p)}}{(r-2)!(rH_{r,2} - H_{r,1}^2)}$. Finally, we

rewrite $(L^T F)_{\mathbf{y},:}$ as $(\alpha(|\mathbf{y}|) - \beta(|\mathbf{y}|))\mathbf{y} + \beta(|\mathbf{y}|)\mathbf{1}$. By the same argument, after the vector multiplication, the second component vanishes and we get $(L^T F (F^T F)^{-1} F \Delta_{\pi\omega})_{\mathbf{y}} = \frac{(F\mathbf{y})\pi - (F\mathbf{y})\omega}{(r-2)!(rH_{r,2} - H_{r,1}^2)}$, which finishes the proof. \square

Lemma 13. *For the score set $\mathcal{F}_{\text{sort}}$, we have $2(r-1)! \|P_{\mathcal{F}_{\text{sort}}} \Delta_{\pi\omega}\|_2^2 = O(r)$. We also have that $\gamma(|\mathbf{y}|)$ defined in Lemma 12 with $|\mathbf{y}| = \lambda r$, $\lambda \in (0, 1)$ vanishes as r approaches infinity: $\gamma(|\mathbf{y}|) = O(\frac{\log^2 r}{r})$. The condition number $\kappa(\mathcal{F}_{\text{sort}})$ grows as $\Theta(\log r)$.*

Proof. To derive an asymptotic bound for $\|P_{\mathcal{F}_{\text{sort}}} \Delta_{\pi\omega}\|_2^2$, we elaborate on the sum of squares $\sum_{p=1}^r \left(\frac{1}{\pi(p)} - \frac{1}{\omega(p)}\right)^2 = 2H_{r,2} - 2\sum_{p=1}^r \frac{1}{\pi(p)\omega(p)} \leq 2H_{r,2}$ and apply the asymptotic bounds for the harmonic numbers $H_{r,1}^2 = \Theta(\log^2 r)$, $H_{r,2} = \Theta(1)$:

$$2(r-1)! \|P_{\mathcal{F}_{\text{sort}}} \Delta_{\pi\omega}\|_2^2 = O\left(\frac{(r)!}{(r-2)!r}\right) = O(r) \quad (71)$$

For the second part of the lemma, we rewrite $\alpha(|\mathbf{y}|)$ and $\beta(|\mathbf{y}|)$:

$$\alpha(|\mathbf{y}|) = \mathfrak{A}_r \left(1 - \lambda\left(1 - \frac{1}{r}\right)\right) - \mathfrak{B}_r \frac{3}{2}(1 - \lambda) - \mathfrak{C}_r \frac{1-\lambda}{r} + o(1) \quad (72)$$

$$\beta(|\mathbf{y}|) = \mathfrak{A}_r (1 - \lambda) - \mathfrak{B}_r \left(1 - \frac{3}{2}\lambda\right) + o(1) \quad (73)$$

$$\alpha(|\mathbf{y}|) - \beta(|\mathbf{y}|) = \mathfrak{A}_r \frac{1}{r} - \mathfrak{B}_r \frac{1}{2} - \mathfrak{C}_r \frac{1-\lambda}{r} + o(1) \quad (74)$$

By definition, we have $\mathfrak{A}_r = \Theta((r-1)! \log r)$, $\mathfrak{B}_r = \Theta((r-2)! \log^2 r)$, $\mathfrak{C}_r = \Theta((r-1)!)$, which gives us

$$\gamma(|\mathbf{y}|) = \frac{\alpha(|\mathbf{y}|) - \beta(|\mathbf{y}|)}{(r-2)!(rH_{r,2} - H_{r,1}^2)} = O\left(\frac{(r-2)! \log^2 r}{(r-1)!}\right) = O\left(\frac{\log^2 r}{r}\right), \quad (75)$$

what was to be shown.

Finally, the asymptotic bound for the condition number of $\mathcal{F}_{\text{sort}}$ trivially follows from its exact expression in Lemma 11 and the asymptotic bounds for the harmonic numbers. \square