
On Structured Prediction Theory with Calibrated Convex Surrogate Losses

Anton Osokin
INRIA/ENS*, Paris, France
HSE[†], Moscow, Russia

Francis Bach
INRIA/ENS*, Paris, France

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

Abstract

We provide novel theoretical insights on structured prediction in the context of *efficient* convex surrogate loss minimization with consistency guarantees. For any task loss, we construct a convex surrogate that can be optimized via stochastic gradient descent and we prove tight bounds on the so-called “calibration function” relating the excess surrogate risk to the actual risk. In contrast to prior related work, we carefully monitor the effect of the exponential number of classes in the learning guarantees as well as on the optimization complexity. As an interesting consequence, we formalize the intuition that some task losses make learning harder than others, and that the classical 0-1 loss is ill-suited for structured prediction.

1 Introduction

Structured prediction is a subfield of machine learning aiming at making multiple interrelated predictions simultaneously. The desired outputs (labels) are typically organized in some structured object such as a sequence, a graph, an image, etc. Tasks of this type appear in many practical domains such as computer vision [34], natural language processing [42] and bioinformatics [19].

The structured prediction setup has at least two typical properties differentiating it from the classical binary classification problems extensively studied in learning theory:

1. Exponential number of classes: this brings both additional computational and statistical challenges. By *exponential*, we mean exponentially large in the size of the natural dimension of output, e.g., the number of all possible sequences is exponential w.r.t. the sequence length.
2. Cost-sensitive learning: in typical applications, prediction mistakes are *not* all equally costly. The prediction error is usually measured with a highly-structured task-specific loss function, e.g., Hamming distance between sequences of multi-label variables or mean average precision for ranking.

Despite many algorithmic advances to tackle structured prediction problems [4, 35], there have been relatively few papers devoted to its theoretical understanding. Notable recent exceptions that made significant progress include Cortes et al. [13] and London et al. [28] (see references therein) which proposed data-dependent generalization error bounds in terms of popular empirical convex surrogate losses such as the structured hinge loss [44, 45, 47]. A question not addressed by these works is whether their algorithms are *consistent*: does minimizing their convex bounds with infinite data lead to the minimization of the task loss as well? Alternatively, the structured probit and ramp losses are consistent [31, 30], but non-convex and thus it is hard to obtain computational guarantees for them. In this paper, we aim at getting the property of consistency for surrogate losses that can be *efficiently* minimized with guarantees, and thus we consider *convex* surrogate losses.

The consistency of convex surrogates is well understood in the case of binary classification [50, 5, 43] and there is significant progress in the case of multi-class 0-1 loss [49, 46] and general multi-

*DI École normale supérieure, CNRS, PSL Research University

[†]National Research University Higher School of Economics

class loss functions [3, 39, 48]. A large body of work specifically focuses on the related tasks of ranking [18, 9, 40] and ordinal regression [37].

Contributions. In this paper, we study consistent convex surrogate losses specifically in the context of an exponential number of classes. We argue that even while being consistent, a convex surrogate might not allow efficient learning. As a concrete example, Ciliberto et al. [10] recently proposed a consistent approach to structured prediction, but the constant in their generalization error bound can be exponentially large as we explain in Section 5. There are two possible sources of difficulties from the optimization perspective: to reach adequate accuracy on the *task* loss, one might need to optimize a surrogate loss to exponentially small accuracy; or to reach adequate accuracy on the *surrogate* loss, one might need an exponential number of algorithm steps because of exponentially large constants in the convergence rate. We propose a theoretical framework that jointly tackles these two aspects and allows to judge the feasibility of efficient learning. In particular, we construct a *calibration function* [43], i.e., a function setting the relationship between accuracy on the surrogate and task losses, and normalize it by the means of convergence rate of an optimization algorithm.

Aiming for the simplest possible application of our framework, we propose a family of convex surrogates that are consistent for any given task loss and can be optimized using stochastic gradient descent. For a special case of our family (quadratic surrogate), we provide a complete analysis including general lower and upper bounds on the calibration function for any task loss, with exact values for the 0-1, block 0-1 and Hamming losses. We observe that to have a tractable learning algorithm, one needs both a structured loss (not the 0-1 loss) and appropriate constraints on the predictor, e.g., in the form of linear constraints for the score vector functions. Our framework also indicates that in some cases it might be beneficial to use non-consistent surrogates. In particular, a non-consistent surrogate might allow optimization only up to specific accuracy, but exponentially faster than a consistent one.

We introduce the structured prediction setting suitable for studying consistency in Sections 2 and 3. We analyze the calibration function for the quadratic surrogate loss in Section 4. We review the related works in Section 5 and conclude in Section 6.

2 Structured prediction setup

In structured prediction, the goal is to predict a structured output $\mathbf{y} \in \mathcal{Y}$ (such as a sequence, a graph, an image) given an input $\mathbf{x} \in \mathcal{X}$. The quality of prediction is measured by a task-dependent *loss function* $L(\hat{\mathbf{y}}, \mathbf{y} \mid \mathbf{x}) \geq 0$ specifying the cost for predicting $\hat{\mathbf{y}}$ when the correct output is \mathbf{y} . In this paper, we consider the case when the number of possible predictions and the number of possible labels are both finite. For simplicity,¹ we also assume that the sets of possible predictions and correct outputs always coincide and do not depend on \mathbf{x} . We refer to this set as the set of labels \mathcal{Y} , denote its cardinality by k , and map its elements to $1, \dots, k$. In this setting, assuming that the loss function depends only on $\hat{\mathbf{y}}$ and \mathbf{y} , but not on \mathbf{x} directly, the loss is defined by a loss matrix $L \in \mathbb{R}^{k \times k}$. We assume that all the elements of the matrix L are non-negative and will use L_{\max} to denote the maximal element. Compared to multi-class classification, k is typically exponentially large in the size of the natural dimension of \mathbf{y} , e.g., contains all possible sequences of symbols from a finite alphabet.

Following standard practices in structured prediction [12, 44], we define the prediction model by a *score function* $\mathbf{f} : \mathcal{X} \rightarrow \mathbb{R}^k$ specifying a score $\mathbf{f}_{\mathbf{y}}(\mathbf{x})$ for each possible output $\mathbf{y} \in \mathcal{Y}$. The final prediction is done by selecting a label with the maximal value of the score

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

with some fixed strategy to resolve ties. To simplify the analysis, we assume that among the labels with maximal scores, the predictor always picks the one with the smallest index.

The goal of prediction-based machine learning consists in finding a predictor that works well on the unseen test set, i.e., data points coming from the same distribution \mathcal{D} as the one generating the training data. One way to formalize this is to minimize the generalization error, often referred to as the actual (or population) *risk* based on the loss L ,

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

Minimizing the actual risk (2) is usually hard. The standard approach is to minimize a *surrogate risk*, which is a different objective easier to optimize, e.g., convex. We define a surrogate loss as a function

¹Our analysis is generalizable to rectangular losses, e.g., ranking losses studied by Ramaswamy et al. [40].

$\Phi : \mathbb{R}^k \times \mathcal{Y} \rightarrow \mathbb{R}$ depending on a score vector $\mathbf{f} = \mathbf{f}(\mathbf{x}) \in \mathbb{R}^k$ and a target label $\mathbf{y} \in \mathcal{Y}$ as input arguments. We denote the \mathbf{y} -th component of \mathbf{f} with $f_{\mathbf{y}}$. The surrogate risk (the Φ -risk) is defined as

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

where the expectation is taken w.r.t. the data-generating distribution \mathcal{D} . To make the minimization of (3) well-defined, we always assume that the surrogate loss Φ is bounded from below and continuous.

Examples of common surrogate losses include the structured hinge-loss [44, 47] $\Phi_{\text{SSVM}}(\mathbf{f}, \mathbf{y}) := \max_{\hat{\mathbf{y}} \in \mathcal{Y}} (f_{\hat{\mathbf{y}}} + L(\hat{\mathbf{y}}, \mathbf{y})) - f_{\mathbf{y}}$, the log loss (maximum likelihood learning) used, e.g., in conditional random fields [25], $\Phi_{\log}(\mathbf{f}, \mathbf{y}) := \log(\sum_{\hat{\mathbf{y}} \in \mathcal{Y}} \exp f_{\hat{\mathbf{y}}}) - f_{\mathbf{y}}$, and their hybrids [38, 21, 22, 41].

In terms of task losses, we consider the unstructured 0-1 loss $L_{01}(\hat{\mathbf{y}}, \mathbf{y}) := [\hat{\mathbf{y}} \neq \mathbf{y}]$,² and the two following structured losses: *block 0-1 loss* with b equal blocks of labels $L_{01,b}(\hat{\mathbf{y}}, \mathbf{y}) := [\hat{\mathbf{y}} \text{ and } \mathbf{y} \text{ are not in the same block}]$; and (normalized) *Hamming loss* between tuples of T binary variables y_t : $L_{\text{Ham},T}(\hat{\mathbf{y}}, \mathbf{y}) := \frac{1}{T} \sum_{t=1}^T [\hat{y}_t \neq y_t]$. To illustrate some aspects of our analysis, we also look at the *mixed loss* $L_{01,b,\eta}$: a convex combination of the 0-1 and block 0-1 losses, defined as $L_{01,b,\eta} := \eta L_{01} + (1 - \eta) L_{01,b}$ for some $\eta \in [0, 1]$.

3 Consistency for structured prediction

3.1 Calibration function

We now formalize the connection between the actual risk \mathcal{R}_L and the surrogate Φ -risk \mathcal{R}_{Φ} via the so-called *calibration function*, see Definition 1 below [5, 49, 43, 18, 3]. As it is standard for this kind of analysis, the setup is *non-parametric*, i.e. it does not take into account the dependency of scores on input variables \mathbf{x} . For now, we assume that a family of score functions $\mathfrak{F}_{\mathcal{F}}$ 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, which will play an important role in our analysis. This setting is equivalent to a pointwise analysis, i.e. looking at the different input \mathbf{x} independently. We bring the dependency on the input back into the analysis in Section 3.3 where we assume a specific family of score functions.

Let $\mathcal{D}_{\mathcal{X}}$ represent the marginal distribution for \mathcal{D} on \mathbf{x} and $\mathbf{P}(\cdot | \mathbf{x})$ denote its conditional given \mathbf{x} . We can now rewrite the risk \mathcal{R}_L and Φ -risk \mathcal{R}_{Φ} as

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

where the conditional risk ℓ and the conditional Φ -risk ϕ depend on a vector of scores \mathbf{f} and a conditional distribution on the set of output labels \mathbf{q} as

$$\ell(\mathbf{f}, \mathbf{q}) := \sum_{c=1}^k q_c L(\text{pred}(\mathbf{f}), c), \quad \phi(\mathbf{f}, \mathbf{q}) := \sum_{c=1}^k q_c \Phi(\mathbf{f}, c).$$

The *calibration function* $H_{\Phi,L,\mathcal{F}}$ between the surrogate loss Φ and the task loss L relates the excess surrogate risk with the actual excess risk via the *excess risk bound*:

$$H_{\Phi,L,\mathcal{F}}(\delta\ell(\mathbf{f}, \mathbf{q})) \leq \delta\phi(\mathbf{f}, \mathbf{q}), \quad \forall \mathbf{f} \in \mathcal{F}, \quad \forall \mathbf{q} \in \Delta_k, \quad (4)$$

where $\delta\phi(\mathbf{f}, \mathbf{q}) = \phi(\mathbf{f}, \mathbf{q}) - \inf_{\hat{\mathbf{f}} \in \mathcal{F}} \phi(\hat{\mathbf{f}}, \mathbf{q})$, $\delta\ell(\mathbf{f}, \mathbf{q}) = \ell(\mathbf{f}, \mathbf{q}) - \inf_{\hat{\mathbf{f}} \in \mathcal{F}} \ell(\hat{\mathbf{f}}, \mathbf{q})$ are the excess risks and Δ_k denotes the probability simplex on k elements.

In other words, to find a vector \mathbf{f} that yields an excess risk smaller than ε , we need to optimize the Φ -risk up to $H_{\Phi,L,\mathcal{F}}(\varepsilon)$ accuracy (in the worst case). We make this statement precise in Theorem 2 below, and now proceed to the formal definition of the calibration function.

Definition 1 (Calibration function). *For a task loss L , a surrogate loss Φ , a set of feasible scores \mathcal{F} , the calibration function $H_{\Phi,L,\mathcal{F}}(\varepsilon)$ (defined for $\varepsilon \geq 0$) equals the infimum excess of the conditional surrogate risk when the excess of the conditional actual risk is at least ε :*

$$H_{\Phi,L,\mathcal{F}}(\varepsilon) := \inf_{\mathbf{f} \in \mathcal{F}, \mathbf{q} \in \Delta_k} \delta\phi(\mathbf{f}, \mathbf{q}) \quad (5)$$

$$\text{s.t.} \quad \delta\ell(\mathbf{f}, \mathbf{q}) \geq \varepsilon. \quad (6)$$

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 $[0, +\infty)$, $H_{\Phi,L,\mathcal{F}}(\varepsilon) \geq 0$, the inequality (4) holds, and $H_{\Phi,L,\mathcal{F}}(0) = 0$. Note that $H_{\Phi,L,\mathcal{F}}$ can be non-convex and even non-continuous (see examples in Figure 1). Also, note that large values of $H_{\Phi,L,\mathcal{F}}(\varepsilon)$ are better.

²Here we use the Iverson bracket notation, i.e., $[A] := 1$ if a logical expression A is true, and zero otherwise.




Figure 1: Calibration functions for the quadratic surrogate Φ_{quad} (12) defined in Section 4 and two different task losses. (a) – the calibration functions for the Hamming loss $L_{\text{Ham},T}$ when used without constraints on the scores, $\mathcal{F} = \mathbb{R}^k$ (in red), and with the tight constraints implying consistency, $\mathcal{F} = \text{span}(L_{\text{Ham},T})$ (in blue). The red curve can grow exponentially slower than the blue one. (b) – the calibration functions for the mixed loss $L_{01,b,\eta}$ with $\eta = 0.4$ (see Section 2 for the definition) when used without constraints on the scores (red) and with tight constraints for the block 0-1 loss (blue). The blue curve represents level-0.2 consistency. The calibration function equals zero for $\varepsilon \leq \eta/2$, but grows exponentially faster than the red curve representing a consistent approach and thus could be better for small η . More details on the calibration functions in this figure are given in Section 4.

3.2 Notion of consistency

We use the calibration function $H_{\Phi,L,\mathcal{F}}$ to set a connection between optimizing the surrogate and task losses by Theorem 2, which is similar to Theorem 3 of Zhang [49].

Theorem 2 (Calibration connection). *Let $H_{\Phi,L,\mathcal{F}}$ be the calibration function between the surrogate loss Φ and the task loss L with feasible set of scores $\mathcal{F} \subseteq \mathbb{R}^k$. Let $\check{H}_{\Phi,L,\mathcal{F}}$ be a convex non-decreasing lower bound of the calibration function. Assume that Φ is continuous and bounded from below. Then, for any $\varepsilon > 0$ with finite $\check{H}_{\Phi,L,\mathcal{F}}(\varepsilon)$ and any $\mathbf{f} \in \mathfrak{F}_{\mathcal{F}}$, we have*

$$\mathcal{R}_{\Phi}(\mathbf{f}) < \mathcal{R}_{\Phi,\mathcal{F}}^* + \check{H}_{\Phi,L,\mathcal{F}}(\varepsilon) \Rightarrow \mathcal{R}_L(\mathbf{f}) < \mathcal{R}_{L,\mathcal{F}}^* + \varepsilon, \quad (7)$$

where $\mathcal{R}_{\Phi,\mathcal{F}}^* := \inf_{\mathbf{f} \in \mathfrak{F}_{\mathcal{F}}} \mathcal{R}_{\Phi}(\mathbf{f})$ and $\mathcal{R}_{L,\mathcal{F}}^* := \inf_{\mathbf{f} \in \mathfrak{F}_{\mathcal{F}}} \mathcal{R}_L(\mathbf{f})$.

Proof. We take the expectation of (4) w.r.t. \mathbf{x} , where the second argument of ℓ is set to the conditional distribution $\mathbf{P}(\cdot | \mathbf{x})$. Then, we apply Jensen’s inequality (since $\check{H}_{\Phi,L,\mathcal{F}}$ is convex) to get

$$\check{H}_{\Phi,L,\mathcal{F}}(\mathcal{R}_L(\mathbf{f}) - \mathcal{R}_{L,\mathcal{F}}^*) \leq \mathcal{R}_{\Phi}(\mathbf{f}) - \mathcal{R}_{\Phi,\mathcal{F}}^* < \check{H}_{\Phi,L,\mathcal{F}}(\varepsilon), \quad (8)$$

which implies (7) by monotonicity of $\check{H}_{\Phi,L,\mathcal{F}}$. \square

A suitable convex non-decreasing lower bound $\check{H}_{\Phi,L,\mathcal{F}}(\varepsilon)$ required by Theorem 2 always exists, e.g., the zero constant. However, in this case Theorem 2 is not informative, because the l.h.s. of (7) is never true. Zhang [49, Proposition 25] claims that $\check{H}_{\Phi,L,\mathcal{F}}$ defined as the lower convex envelope of the calibration function $H_{\Phi,L,\mathcal{F}}$ satisfies $\check{H}_{\Phi,L,\mathcal{F}}(\varepsilon) > 0, \forall \varepsilon > 0$, if $H_{\Phi,L,\mathcal{F}}(\varepsilon) > 0, \forall \varepsilon > 0$, and, e.g., the set of labels is finite. This statement implies that an informative $\check{H}_{\Phi,L,\mathcal{F}}$ always exists and allows to characterize consistency through properties of the calibration function $H_{\Phi,L,\mathcal{F}}$.

We now define a notion of *level- η consistency*, which is more general than consistency.

Definition 3 (level- η consistency). *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.*

Looking solely at (standard level-0) consistency vs. inconsistency might be too coarse to capture practical properties related to optimization accuracy (see, e.g., [29]). For example, if $H_{\Phi,L,\mathcal{F}}(\varepsilon) = 0$ only for very small values of ε , then the method can still optimize the actual risk up to a certain level which might be good enough in practice, especially if it means that it can be optimized faster. Examples of calibration functions for consistent and inconsistent surrogate losses are shown in Figure 1.

Other notions of consistency. Definition 3 with $\eta = 0$ and $\mathcal{F} = \mathbb{R}^k$ results in the standard setting often appearing in the literature. In particular, in this case Theorem 2 implies Fisher consistency as

formulated, e.g., by Pedregosa et al. [37] for general losses and Lin [27] for binary classification. This setting is also closely related to many definitions of consistency used in the literature. For example, for a bounded from below and continuous surrogate, it is equivalent to infinite-sample consistency [49], classification calibration [46], edge-consistency [18], (L, \mathbb{R}^k) -calibration [39], prediction calibration [48]. See [49, Appendix A] for the detailed discussion.

Role of \mathcal{F} . Let the *approximation error* for the restricted set of scores \mathcal{F} be defined as $\mathcal{R}_{L,\mathcal{F}}^* - \mathcal{R}_L^* := \inf_{f \in \mathcal{F}} \mathcal{R}_L(f) - \inf_f \mathcal{R}_L(f)$. For any conditional distribution q , the score vector $\mathbf{f} := -L\mathbf{q}$ will yield an optimal prediction. Thus the condition $\text{span}(L) \subseteq \mathcal{F}$ is sufficient for \mathcal{F} to have zero approximation error for any distribution \mathcal{D} , and for our 0-consistency condition to imply the standard Fisher consistency with respect to L . In the following, we will see that a restricted \mathcal{F} can both play a role for computational efficiency as well as statistical efficiency (thus losses with smaller $\text{span}(L)$ might be easier to work with).

3.3 Connection to optimization accuracy and statistical efficiency

The scale of a calibration function is not intrinsically well-defined: we could multiply the surrogate function by a scalar and it would multiply the calibration function by the same scalar, without changing the optimization problem. Intuitively, we would like the surrogate loss to be of order 1. If with this scale the calibration function is exponentially small (has a $1/k$ factor), then we have strong evidence that the stochastic optimization will be difficult (and thus learning will be slow).

To formalize this intuition, we add to the picture the *complexity* of optimizing the surrogate loss with a *stochastic approximation* algorithm. By using a scale-invariant convergence rate, we provide a natural normalization of the calibration function. The following two observations are central to the theoretical insights provided in our work:

1. **Scale.** For a properly scaled surrogate loss, the *scale* of the calibration function is a good indication of whether a stochastic approximation algorithm will take a large number of iterations (in the worst case) to obtain guarantees of small excess of the actual risk (and vice-versa, a large coefficient indicates a small number of iterations). The actual verification requires computing the normalization quantities given in Theorem 6 below.
2. **Statistics.** The bound on the number of iterations directly relates to the number of training examples that would be needed to learn, if we see each iteration of the stochastic approximation algorithm as using one training example to optimize the expected surrogate.

To analyze the statistical convergence of surrogate risk optimization, we have to specify the set of score functions that we work with. We assume that the structure on input $\mathbf{x} \in \mathcal{X}$ is defined by a positive definite kernel $K : \mathcal{X} \times \mathcal{X} \rightarrow \mathbb{R}$. We denote the corresponding reproducing kernel Hilbert space (RKHS) by \mathcal{H} and its explicit feature map by $\psi(\mathbf{x}) \in \mathcal{H}$. By the reproducing property, we have $\langle f, \psi(\mathbf{x}) \rangle_{\mathcal{H}} = f(\mathbf{x})$ for all $\mathbf{x} \in \mathcal{X}$, $f \in \mathcal{H}$, where $\langle \cdot, \cdot \rangle_{\mathcal{H}}$ is the inner product in the RKHS. We define the subspace of allowed scores $\mathcal{F} \subseteq \mathbb{R}^k$ via the span of the columns of a matrix $F \in \mathbb{R}^{k \times r}$. The matrix F explicitly defines the structure of the score function. With this notation, we will assume that the score function is of the form $f(\mathbf{x}) = FW\psi(\mathbf{x})$, where $W : \mathcal{H} \rightarrow \mathbb{R}^r$ is a linear operator to be learned (a matrix if \mathcal{H} is of finite dimension) that represents a collection of r elements in \mathcal{H} , transforming $\psi(\mathbf{x})$ to a vector in \mathbb{R}^r by applying the RKHS inner product r times.³ Note that for structured losses, we usually have $r \ll k$. The set of all score functions is thus obtained by varying W in this definition and is denoted by $\mathfrak{F}_{F,\mathcal{H}}$. As a concrete example of a score family $\mathfrak{F}_{F,\mathcal{H}}$ for structured prediction, consider the standard sequence model with unary and pairwise potentials. In this case, the dimension r equals $Ts + (T-1)s^2$, where T is the sequence length and s is the number of labels of each variable. The columns of the matrix F consist of $2T-1$ groups (one for each unary and pairwise potential). Each row of F has exactly one entry equal to one in each column group (with zeros elsewhere).

In this setting, we use the online projected averaged stochastic subgradient descent ASGD⁴ (stochastic w.r.t. data $(\mathbf{x}^{(n)}, \mathbf{y}^{(n)}) \sim \mathcal{D}$) to minimize the surrogate risk directly [6]. The n -th update consists in

$$W^{(n)} := P_D [W^{(n-1)} - \gamma^{(n)} F^\top \nabla \Phi \psi(\mathbf{x}^{(n)})^\top], \quad (9)$$

³Note that if $\text{rank}(F) = r$, our setup is equivalent to assuming a *joint kernel* [47] in the product form: $K_{\text{joint}}((\mathbf{x}, c), (\mathbf{x}', c')) := K(\mathbf{x}, \mathbf{x}') F(c, :) F(c', :)^T$, where $F(c, :)$ is the row c for matrix F .

⁴See, e.g., [36] for the formal setup of kernel ASGD.

where $F^\top \nabla \Phi \psi(\mathbf{x}^{(n)})^\top : \mathcal{H} \rightarrow \mathbb{R}^r$ is the stochastic functional gradient, $\gamma^{(n)}$ is the step size and P_D is the projection on the ball of radius D w.r.t. the Hilbert–Schmidt norm⁵. The vector $\nabla \Phi \in \mathbb{R}^k$ is a regular gradient of the sampled surrogate $\Phi(\mathbf{f}(\mathbf{x}^{(n)}), \mathbf{y}^{(n)})$ w.r.t. the scores, $\nabla \Phi = \nabla_{\mathbf{f}} \Phi(\mathbf{f}, \mathbf{y}^{(n)})|_{\mathbf{f}=\mathbf{f}(\mathbf{x}^{(n)})}$. We wrote the above update using an explicit feature map ψ for notational simplicity, but kernel ASGD can also be implemented without it by using the kernel trick. The convergence properties of ASGD in RKHS are analogous to the finite-dimensional ASGD because they rely on dimension-free quantities. To use a simple convergence analysis, we follow Ciliberto et al. [10] and make the following simplifying assumption:

Assumption 4 (Well-specified optimization w.r.t. the function class $\mathfrak{F}_{F, \mathcal{H}}$). *The distribution \mathcal{D} is such that $\mathcal{R}_{\Phi, \mathcal{F}}^* := \inf_{\mathbf{f} \in \mathfrak{F}_{\mathcal{F}}} \mathcal{R}_{\Phi}(\mathbf{f})$ has some global minimum \mathbf{f}^* that also belongs to $\mathfrak{F}_{F, \mathcal{H}}$.*

Assumption 4 simply means that each row of W^* defining \mathbf{f}^* belongs to the RKHS \mathcal{H} implying a finite norm $\|W^*\|_{HS}$. Assumption 4 can be relaxed if the kernel K is universal, but then the convergence analysis becomes much more complicated [36].

Theorem 5 (Convergence rate). *Under Assumption 4 and assuming that (i) the functions $\Phi(\mathbf{f}, \mathbf{y})$ are bounded from below and convex w.r.t. $\mathbf{f} \in \mathbb{R}^k$ for all $\mathbf{y} \in \mathcal{Y}$; (ii) the expected square of the norm of the stochastic gradient is bounded, $\mathbb{E}_{(\mathbf{x}, \mathbf{y}) \sim \mathcal{D}} \|F^\top \nabla \Phi \psi(\mathbf{x})^\top\|_{HS}^2 \leq M^2$ and (iii) $\|W^*\|_{HS} \leq D$, then running the ASGD algorithm (9) with the constant step-size $\gamma := \frac{2D}{M\sqrt{N}}$ for N steps admits the following expected suboptimality for the averaged iterate $\bar{\mathbf{f}}^{(N)}$:*

$$\mathbb{E}[\mathcal{R}_{\Phi}(\bar{\mathbf{f}}^{(N)})] - \mathcal{R}_{\Phi, \mathcal{F}}^* \leq \frac{2DM}{\sqrt{N}} \quad \text{where} \quad \bar{\mathbf{f}}^{(N)} := \frac{1}{N} \sum_{n=1}^N F W^{(n)} \psi(\mathbf{x}^{(n)})^\top. \quad (10)$$

Theorem 5 is a straight-forward extension of classical results [33, 36].

By combining the convergence rate of Theorem 5 with Theorem 2 that connects the surrogate and actual risks, we get Theorem 6 which explicitly gives the number of iterations required to achieve ε accuracy on the expected population risk (see App. A for the proof). Note that since ASGD is applied in an online fashion, Theorem 6 also serves as the sample complexity bound, i.e., says how many samples are needed to achieve ε target accuracy (compared to the best prediction rule if \mathcal{F} has zero approximation error).

Theorem 6 (Learning complexity). *Under the assumptions of Theorem 5, for any $\varepsilon > 0$, the random (w.r.t. the observed training set) output $\bar{\mathbf{f}}^{(N)} \in \mathfrak{F}_{F, \mathcal{H}}$ of the ASGD algorithm after*

$$N > N^* := \frac{4D^2 M^2}{\hat{H}_{\Phi, L, \mathcal{F}}^2(\varepsilon)} \quad (11)$$

iterations has the expected excess risk bounded with ε , i.e., $\mathbb{E}[\mathcal{R}_L(\bar{\mathbf{f}}^{(N)})] < \mathcal{R}_{L, \mathcal{F}}^ + \varepsilon$.*

4 Calibration function analysis for quadratic surrogate

A major challenge to applying Theorem 6 is the computation of the calibration function $H_{\Phi, L, \mathcal{F}}$. In App. C, we present a generalization to arbitrary multi-class losses of a surrogate loss class from Zhang [49, Section 4.4.2] that is consistent for any task loss L . Here, we consider the simplest example of this family, called the *quadratic surrogate* Φ_{quad} , which has the advantage that we can bound or even compute exactly its calibration function. We define the quadratic surrogate as

$$\Phi_{\text{quad}}(\mathbf{f}, \mathbf{y}) := \frac{1}{2k} \|\mathbf{f} + L(\cdot, \mathbf{y})\|_2^2 = \frac{1}{2k} \sum_{c=1}^k (f_c^2 + 2f_c L(c, \mathbf{y}) + L(c, \mathbf{y})^2). \quad (12)$$

One simple sufficient condition for the surrogate (12) to be consistent and also to have zero approximation error is that \mathcal{F} fully contains $\text{span}(L)$. To make the dependence on the score subspace explicit, we parameterize it with a matrix $F \in \mathbb{R}^{k \times r}$ with the number of columns r typically being much smaller than the number of labels k . With this notation, we have $\mathcal{F} = \text{span}(F) = \{F\boldsymbol{\theta} \mid \boldsymbol{\theta} \in \mathbb{R}^r\}$, and the dimensionality of \mathcal{F} equals the rank of F , which is at most r .⁶

⁵The Hilbert–Schmidt norm of a linear operator A is defined as $\|A\|_{HS} = \sqrt{\text{tr} A^\dagger A}$ where A^\dagger is the adjoint operator. In the case of finite dimension, the Hilbert–Schmidt norm coincides with the Frobenius matrix norm.

⁶Evaluating Φ_{quad} requires computing $F^\top F$ and $F^\top L(\cdot, \mathbf{y})$ for which direct computation is intractable when k is exponential, but which can be done in closed form for the structured losses we consider (the Hamming and block 0-1 loss). More generally, these operations require suitable inference algorithms. See also App. F.

For the quadratic surrogate (12), the excess of the expected surrogate takes a simple form:

$$\delta\phi_{\text{quad}}(F\theta, \mathbf{q}) = \frac{1}{2k} \|F\theta + L\mathbf{q}\|_2^2. \quad (13)$$

Equation (13) holds under the assumption that the subspace \mathcal{F} contains the column space of the loss matrix $\text{span}(L)$, which also means that the set \mathcal{F} contains the optimal prediction for any \mathbf{q} (see Lemma 9 in App. B for the proof). Importantly, the function $\delta\phi_{\text{quad}}(F\theta, \mathbf{q})$ is jointly convex in the conditional probability \mathbf{q} and parameters θ , which simplifies its analysis.

Lower bound on the calibration function. We now present our main technical result: a lower bound on the calibration function for the surrogate loss Φ_{quad} (12). This lower bound characterizes the easiness of learning with this surrogate given the scaling intuition mentioned in Section 3.3. The proof of Theorem 7 is given in App. D.1.

Theorem 7 (Lower bound on $H_{\Phi_{\text{quad}}}$). *For any task loss L , its quadratic surrogate Φ_{quad} , and a score subspace \mathcal{F} containing the column space of L , the calibration function can be lower bounded:*

$$H_{\Phi_{\text{quad}}, L, \mathcal{F}}(\varepsilon) \geq \frac{\varepsilon^2}{2k \max_{i \neq j} \|P_{\mathcal{F}} \Delta_{ij}\|_2^2} \geq \frac{\varepsilon^2}{4k}, \quad (14)$$

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 .

Lower bound for specific losses. We now discuss the meaning of the bound (14) for some specific losses (the detailed derivations are given in App. D.3). For the 0-1, block 0-1 and Hamming losses (L_{01} , $L_{01,b}$ and $L_{\text{Ham}, T}$, respectively) with the smallest possible score subspaces \mathcal{F} , the bound (14) gives $\frac{\varepsilon^2}{4k}$, $\frac{\varepsilon^2}{4b}$ and $\frac{\varepsilon^2}{8T}$, respectively. All these bounds are tight (see App. E). However, if $\mathcal{F} = \mathbb{R}^k$ the bound (14) is not tight for the block 0-1 and mixed losses (see also App. E). In particular, the bound (14) cannot detect level- η consistency for $\eta > 0$ (see Def. 3) and does not change when the loss changes, but the score subspace stays the same.

Upper bound on the calibration function. Theorem 8 below gives an upper bound on the calibration function holding for unconstrained scores, i.e., $\mathcal{F} = \mathbb{R}^k$ (see the proof in App. D.2). This result shows that without some appropriate constraints on the scores, efficient learning is not guaranteed (in the worst case) because of the $1/k$ scaling of the calibration function.

Theorem 8 (Upper bound on $H_{\Phi_{\text{quad}}}$). *If a loss matrix L with $L_{\max} > 0$ defines a pseudometric⁷ on labels and there are no constraints on the scores, i.e., $\mathcal{F} = \mathbb{R}^k$, then the calibration function for the quadratic surrogate Φ_{quad} can be upper bounded: $H_{\Phi_{\text{quad}}, L, \mathbb{R}^k}(\varepsilon) \leq \frac{\varepsilon^2}{2k}$, $0 \leq \varepsilon \leq L_{\max}$.*

From our lower bound in Theorem 7 (which guarantees consistency), the natural constraint on the score is $\mathcal{F} = \text{span}(L)$, with the dimension of this space giving an indication of the intrinsic “difficulty” of a loss. Computations for the lower bounds in some specific cases (see App. D.3 for details) show that the 0-1 loss is “hard” while the block 0-1 loss and the Hamming loss are “easy”. Note that in all these cases the lower bound (14) is tight, see the discussion below.

Exact calibration functions. Note that the bounds proven in Theorems 7 and 8 imply that, in the case of no constraints on the scores $\mathcal{F} = \mathbb{R}^k$, for the 0-1, block 0-1 and Hamming losses, we have

$$\frac{\varepsilon^2}{4k} \leq H_{\Phi_{\text{quad}}, L, \mathbb{R}^k}(\varepsilon) \leq \frac{\varepsilon^2}{2k}, \quad (15)$$

where L is the matrix defining a loss. For completeness, in App. E, we compute the exact calibration functions for the 0-1 and block 0-1 losses. Note that the calibration function for the **0-1 loss** equals the lower bound, illustrating the worst-case scenario. To get some intuition, an example of a conditional distribution \mathbf{q} that gives the (worst case) value to the calibration function (for several losses) is $q_i = \frac{1}{2} + \frac{\varepsilon}{2}$, $q_j = \frac{1}{2} - \frac{\varepsilon}{2}$ and $q_c = 0$ for $c \notin \{i, j\}$. See the proof of Proposition 12 in App. E.1.

In what follows, we provide the calibration functions in the cases with constraints on the scores. For the **block 0-1 loss** with b equal blocks and under constraints that the scores within blocks are equal, the calibration function equals (see Proposition 14 of App. E.2)

$$H_{\Phi_{\text{quad}}, L_{01,b}, \mathcal{F}_{01,b}}(\varepsilon) = \frac{\varepsilon^2}{4b}, \quad 0 \leq \varepsilon \leq 1. \quad (16)$$

⁷A pseudometric is a function $d(a, b)$ satisfying the following axioms: $d(x, y) \geq 0$, $d(x, x) = 0$ (but possibly $d(x, y) = 0$ for some $x \neq y$), $d(x, y) = d(y, x)$, $d(x, z) \leq d(x, y) + d(y, z)$.

For the **Hamming loss** defined over T binary variables and under constraints implying separable scores, the calibration function equals (see Proposition 15 in App. E.3)

$$H_{\Phi_{\text{quad}}, L_{\text{Ham}}, T, \mathcal{F}_{\text{Ham}}, T}(\varepsilon) = \frac{\varepsilon^2}{8T}, \quad 0 \leq \varepsilon \leq 1. \quad (17)$$

The calibration functions (16) and (17) depend on the quantities representing the actual complexities of the loss (the number of blocks b and the length of the sequence T) and can be exponentially larger than the upper bound for the unconstrained case.

In the case of **mixed 0-1 and block 0-1 loss**, if the scores \mathbf{f} are constrained to be equal inside the blocks, i.e., belong to the subspace $\mathcal{F}_{01,b} = \text{span}(L_{01,b}) \subsetneq \mathbb{R}^k$, then the calibration function is equal to 0 for $\varepsilon \leq \frac{\eta}{2}$, implying inconsistency (and also note that the approximation error can be as big as η for $\mathcal{F}_{01,b}$). However, for $\varepsilon > \frac{\eta}{2}$, the calibration function is of the order $\frac{1}{b}(\varepsilon - \frac{\eta}{2})^2$. See Figure 1b for the illustration of this calibration function and Proposition 17 of App. E.4 for the exact formulation and the proof. Note that while the calibration function for the constrained case is inconsistent, its value can be exponentially larger than the one for the unconstrained case for ε big enough and when the blocks are exponentially large (see Proposition 16 of App. E.4).

Computation of the SGD constants. Applying the learning complexity Theorem 6 requires to compute the quantity DM where D bounds the norm of the optimal solution and M bounds the expected square of the norm of the stochastic gradient. In App. F, we provide a way to bound this quantity for our quadratic surrogate (12) under the simplifying assumption that each conditional $q_c(\mathbf{x})$ (seen as function of \mathbf{x}) belongs to the RKHS \mathcal{H} (which implies Assumption 4). In particular, we get

$$DM = L_{\max}^2 \xi(\kappa(F) \sqrt{r} R Q_{\max}), \quad \xi(z) = z^2 + z, \quad (18)$$

where $\kappa(F)$ is the condition number of the matrix F , R is an upper bound on the RKHS norm of object feature maps $\|\psi(\mathbf{x})\|_{\mathcal{H}}$. We define Q_{\max} as an upper bound on $\sum_{c=1}^k \|q_c\|_{\mathcal{H}}$ (can be seen as the generalization of the inequality $\sum_{c=1}^k q_c \leq 1$ for probabilities). The constants R and Q_{\max} depend on the data, the constant L_{\max} depends on the loss, r and $\kappa(F)$ depend on the choice of matrix F .

We compute the constant DM for the specific losses that we considered in App. F.1. For the 0-1, block 0-1 and Hamming losses, we have $DM = O(k)$, $DM = O(b)$ and $DM = O(\log_2^3 k)$, respectively. These computations indicate that the quadratic surrogate allows efficient learning for structured block 0-1 and Hamming losses, but that the convergence could be slow in the worst case for the 0-1 loss.

5 Related works

Consistency for multi-class problems. Building on significant progress for the case of binary classification, see, e.g. [5], there has been a lot of interest in the multi-class case. Zhang [49] and Tewari & Bartlett [46] analyze the consistency of many existing surrogates for the 0-1 loss. Gao & Zhou [20] focus on multi-label classification. Narasimhan et al. [32] provide a consistent algorithm for arbitrary multi-class loss defined by a function of the confusion matrix. Recently, Ramaswamy & Agarwal [39] introduce the notion of convex calibrated dimension, as the minimal dimensionality of the score vector that is required for consistency. In particular, they showed that for the Hamming loss on T binary variables, this dimension is at most T . In our analysis, we use scores of rank $(T+1)$, see (35) in App. D.3, yielding a similar result.

The task of ranking has attracted a lot of attention and [18, 8, 9, 40] analyze different families of surrogate and task losses proving their (in-)consistency. In this line of work, Ramaswamy et al. [40] propose a quadratic surrogate for an arbitrary low rank loss which is related to our quadratic surrogate (12). They also prove that several important ranking losses, i.e., precision@q, expected rank utility, mean average precision and pairwise disagreement, are of low-rank. We conjecture that our approach is compatible with these losses and leave precise connections as future work.

Structured SVM (SSVM) and friends. SSVM [44, 45, 47] is one of the most used convex surrogates for tasks with structured outputs, thus, its consistency has been a question of great interest. It is known that Crammer-Singer multi-class SVM [15], which SSVM is built on, is not consistent for 0-1 loss unless there is a majority class with probability at least $\frac{1}{2}$ [49, 31]. However, it is consistent for the “abstain” and ordinal losses in the case of 3 classes [39]. Structured ramp loss and probit surrogates are closely related to SSVM and are consistent [31, 16, 30, 23], but not convex.

Recently, Doğan et al. [17] categorized different versions of multi-class SVM and analyzed them from Fisher and universal consistency point of views. In particular, they highlight differences between Fisher and universal consistency and give examples of surrogates that are Fisher consistent, but not universally consistent and vice versa. They also highlight that the Crammer-Singer SVM is neither Fisher, not universally consistent even with a careful choice of regularizer.

Quadratic surrogates for structured prediction. Ciliberto et al. [10] and Brouard et al. [7] consider minimizing $\sum_{i=1}^n \|g(x_i) - \psi_o(y_i)\|_{\mathcal{H}}^2$ aiming to match the RKHS embedding of inputs $g : \mathcal{X} \rightarrow \mathcal{H}$ to the feature maps of outputs $\psi_o : \mathcal{Y} \rightarrow \mathcal{H}$. In their frameworks, the task loss is not considered at the learning stage, but only at the prediction stage. Our quadratic surrogate (12) depends on the loss directly. The empirical risk defined by both their and our objectives can be minimized analytically with the help of the kernel trick and, moreover, the resulting predictors are identical. However, performing such computation in the case of large dataset can be intractable and the generalization properties have to be taken care of, e.g., by the means of regularization. In the large-scale scenario, it is more natural to apply stochastic optimization (e.g., kernel ASGD) that directly minimizes the population risk and has better dependency on the dataset size. When combined with stochastic optimization, the two approaches lead to different behavior. In our framework, we need to estimate $r = \text{rank}(L)$ scalar functions, but the alternative needs to estimate k functions (if, e.g., $\psi_o(y) = \mathbf{e}_y \in \mathbb{R}^k$), which results in significant differences for low-rank losses, such as block 0-1 and Hamming.

Calibration functions. Bartlett et al. [5] and Steinwart [43] provide calibration functions for most existing surrogates for binary classification. All these functions differ in term of shape, but are roughly similar in terms of constants. Pedregosa et al. [37] generalize these results to the case of ordinal regression. However, their calibration functions have at best a $1/k$ factor if the surrogate is normalized w.r.t. the number of classes. The task of ranking has been of significant interest. However, most of the literature [e.g., 11, 14, 24, 1], only focuses on calibration functions (in the form of regret bounds) for bipartite ranking, which is more akin to cost-sensitive binary classification.

Ávila Pires et al. [3] generalize the theoretical framework developed by Steinwart [43] and present results for the multi-class SVM of Lee et al. [26] (the score vectors are constrained to sum to zero) that can be built for any task loss of interest. Their surrogate Φ is of the form $\sum_{c \in \mathcal{Y}} L(c, y) a(f_c)$ where $\sum_{c \in \mathcal{Y}} f_c = 0$ and $a(f)$ is some convex function with all subgradients at zero being positive. The recent work by Ávila Pires & Szepesvári [2] refines the results, but specifically for the case of 0-1 loss. In this line of work, the surrogate is typically not normalized by k , and if normalized the calibration functions have the constant $1/k$ appearing.

Finally, Ciliberto et al. [10] provide the calibration function for their quadratic surrogate. Assuming that the loss can be represented as $L(\hat{y}, y) = \langle V\psi_o(\hat{y}), \psi_o(y) \rangle_{\mathcal{H}_Y}$, $\hat{y}, y \in \mathcal{Y}$ (this assumption can always be satisfied in the case of a finite number of labels, by taking V as the loss matrix L and $\psi_o(y) := \mathbf{e}_y \in \mathbb{R}^k$ where \mathbf{e}_y is the y -th vector of the standard basis in \mathbb{R}^k). In their Theorem 2, they provide an excess risk bound leading to a lower bound on the corresponding calibration function $H_{\Phi, L, \mathbb{R}^k}(\varepsilon) \geq \frac{\varepsilon^2}{c_{\Delta}^2}$ where a constant $c_{\Delta} = \|V\|_2 \max_{y \in \mathcal{Y}} \|\psi_o(y)\|$ simply equals the spectral norm of the loss matrix for the finite-dimensional construction provided above. However, the spectral norm of the loss matrix is exponentially large even for highly structured losses such as the block 0-1 and Hamming losses, i.e., $\|L_{01,b}\|_2 = k - \frac{k}{b}$, $\|L_{\text{Ham}, T}\|_2 = \frac{k}{2}$. This conclusion puts the objective of Ciliberto et al. [10] in line with ours when no constraints are put on the scores.

6 Conclusion

In this paper, we studied the consistency of convex surrogate losses specifically in the context of structured prediction. We analyzed calibration functions and proposed an optimization-based normalization aiming to connect consistency with the existence of efficient learning algorithms. Finally, we instantiated all components of our framework for several losses by computing the calibration functions and the constants coming from the normalization. By carefully monitoring exponential constants, we highlighted the difference between tractable and intractable task losses.

These were first steps in advancing our theoretical understanding of consistent structured prediction. Further steps include analyzing more losses such as the low-rank ranking losses studied by Ramaswamy et al. [40] and, instead of considering constraints on the scores, one could instead put constraints on the set of distributions to investigate the effect on the calibration function.

Acknowledgements

We would like to thank Pascal Germain for useful discussions. This work was partly supported by the ERC grant Activia (no. 307574), the NSERC Discovery Grant RGPIN-2017-06936 and the MSR-INRIA Joint Center.

References

- [1] Agarwal, Shivani. Surrogate regret bounds for bipartite ranking via strongly proper losses. *Journal of Machine Learning Research (JMLR)*, 15(1):1653–1674, 2014.
- [2] Ávila Pires, Bernardo and Szepesvári, Csaba. Multiclass classification calibration functions. *arXiv*, 1609.06385v1, 2016.
- [3] Ávila Pires, Bernardo, Ghavamzadeh, Mohammad, and Szepesvári, Csaba. Cost-sensitive multiclass classification risk bounds. In *ICML*, 2013.
- [4] Bakir, Gökhhan, Hofmann, Thomas, Schölkopf, Bernhard, Smola, Alexander J., Taskar, Ben, and Vishwanathan, S.V.N. *Predicting Structured Data*. MIT press, 2007.
- [5] 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.
- [6] Bousquet, Olivier and Bottou, Léon. The tradeoffs of large scale learning. In *NIPS*, 2008.
- [7] Brouard, Céline, Szafranski, Marie, and d’Alché-Buc, Florence. Input output kernel regression: Supervised and semi-supervised structured output prediction with operator-valued kernels. *Journal of Machine Learning Research (JMLR)*, 17(176):1–48, 2016.
- [8] Buffoni, David, Gallinari, Patrick, Usunier, Nicolas, and Calauzènes, Clément. Learning scoring functions with order-preserving losses and standardized supervision. In *ICML*, 2011.
- [9] Calauzènes, Clément, Usunier, Nicolas, and Gallinari, Patrick. On the (non-)existence of convex, calibrated surrogate losses for ranking. In *NIPS*, 2012.
- [10] Ciliberto, Carlo, Rudi, Alessandro, and Rosasco, Lorenzo. A consistent regularization approach for structured prediction. In *NIPS*, 2016.
- [11] Cléménçon, Stéphan, Lugosi, Gábor, and Vayatis, Nicolas. Ranking and empirical minimization of U-statistics. *The Annals of Statistics*, pp. 844–874, 2008.
- [12] Collins, Michael. Discriminative training methods for hidden Markov models: Theory and experiments with perceptron algorithms. In *EMNLP*, 2002.
- [13] Cortes, Corinna, Kuznetsov, Vitaly, Mohri, Mehryar, and Yang, Scott. Structured prediction theory based on factor graph complexity. In *NIPS*, 2016.
- [14] Cossock, David and Zhang, Tong. Statistical analysis of bayes optimal subset ranking. *IEEE Transactions on Information Theory*, 54(11):5140–5154, 2008.
- [15] 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.
- [16] Do, Chuong B., Le, Quoc, Teo, Choon Hui, Chapelle, Olivier, and Smola, Alex. Tighter bounds for structured estimation. In *NIPS*, 2009.
- [17] Doğan, Ürün, Glasmachers, Tobias, and Igel, Christian. A unified view on multi-class support vector classification. *Journal of Machine Learning Research (JMLR)*, 17(45):1–32, 2016.
- [18] Duchi, John C., Mackey, Lester W., and Jordan, Michael I. On the consistency of ranking algorithms. In *ICML*, 2010.
- [19] Durbin, Richard, Eddy, Sean, Krogh, Anders, and Mitchison, Graeme. *Biological sequence analysis: probabilistic models of proteins and nucleic acids*. Cambridge university press, 1998.

- [20] Gao, Wei and Zhou, Zhi-Hua. On the consistency of multi-label learning. In *COLT*, 2011.
- [21] Gimpel, Kevin and Smith, Noah A. Softmax-margin CRFs: Training loglinear models with cost functions. In *NAACL*, 2010.
- [22] Hazan, Tamir and Urtasun, Raquel. A primal-dual message-passing algorithm for approximated large scale structured prediction. In *NIPS*, 2010.
- [23] Keshet, Joseph. Optimizing the measure of performance in structured prediction. In *Advanced Structured Prediction*. MIT Press, 2014.
- [24] Kotlowski, Wojciech, Dembczynski, Krzysztof, and Huellermeier, Eyke. Bipartite ranking through minimization of univariate loss. In *ICML*, 2011.
- [25] Lafferty, John, McCallum, Andrew, and Pereira, Fernando. Conditional random fields: Probabilistic models for segmenting and labeling sequence data. In *ICML*, 2001.
- [26] Lee, Yoonkyung, Lin, Yi, and Wahba, Grace. Multicategory 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.
- [27] Lin, Yi. A note on margin-based loss functions in classification. *Statistics & Probability Letters*, 68(1):73–82, 2004.
- [28] London, Ben, Huang, Bert, and Getoor, Lise. Stability and generalization in structured prediction. *Journal of Machine Learning Research (JMLR)*, 17(222):1–52, 2016.
- [29] Long, Phil and Servedio, Rocco. Consistency versus realizable H-consistency for multiclass classification. In *ICML*, 2013.
- [30] McAllester, D. A. and Keshet, J. Generalization bounds and consistency for latent structural probit and ramp loss. In *NIPS*, 2011.
- [31] McAllester, David. Generalization bounds and consistency for structured labeling. In *Predicting Structured Data*. MIT Press, 2007.
- [32] Narasimhan, Harikrishna, Ramaswamy, Harish G., Saha, Aadirupa, and Agarwal, Shivani. Consistent multiclass algorithms for complex performance measures. In *ICML*, 2015.
- [33] 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.
- [34] Nowozin, Sebastian and Lampert, Christoph H. Structured learning and prediction in computer vision. *Foundations and Trends in Computer Graphics and Vision*, 6(3–4):185–365, 2011.
- [35] Nowozin, Sebastian, Gehler, Peter V., Jancsary, Jeremy, and Lampert, Christoph H. *Advanced Structured Prediction*. MIT Press, 2014.
- [36] Orabona, Francesco. Simultaneous model selection and optimization through parameter-free stochastic learning. In *NIPS*, 2014.
- [37] 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.
- [38] Pletscher, Patrick, Ong, Cheng Soon, and Buhmann, Joachim M. Entropy and margin maximization for structured output learning. In *ECML PKDD*, 2010.
- [39] Ramaswamy, Harish G. and Agarwal, Shivani. Convex calibration dimension for multiclass loss matrices. *Journal of Machine Learning Research (JMLR)*, 17(14):1–45, 2016.
- [40] 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.
- [41] Shi, Qinfeng, Reid, Mark, Caetano, Tiberio, van den Hengel, Anton, and Wang, Zhenhua. A hybrid loss for multiclass and structured prediction. *IEEE transactions on pattern analysis and machine intelligence (TPAMI)*, 37(1):2–12, 2015.

- [42] Smith, Noah A. Linguistic structure prediction. *Synthesis lectures on human language technologies*, 4(2):1–274, 2011.
- [43] Steinwart, Ingo. How to compare different loss functions and their risks. *Constructive Approximation*, 26(2):225–287, 2007.
- [44] Taskar, Ben, Guestrin, Carlos, and Koller, Daphne. Max-margin markov networks. In *NIPS*, 2003.
- [45] Taskar, Ben, Chatalbashev, Vassil, Koller, Daphne, and Guestrin, Carlos. Learning structured prediction models: a large margin approach. In *ICML*, 2005.
- [46] Tewari, Ambuj and Bartlett, Peter L. On the consistency of multiclass classification methods. *Journal of Machine Learning Research (JMLR)*, 8:1007–1025, 2007.
- [47] 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.
- [48] Williamson, Robert C., Vernet, Elodie, and Reid, Mark D. Composite multiclass losses. *Journal of Machine Learning Research (JMLR)*, 17(223):1–52, 2016.
- [49] Zhang, Tong. Statistical analysis of some multi-category large margin classification methods. *Journal of Machine Learning Research (JMLR)*, 5:1225–1251, 2004.
- [50] 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)

On Structured Prediction Theory with Calibrated Convex Surrogate Losses

Outline

Section A: Proof of learning complexity Theorem 6.
 Section B: Technical lemmas useful for the proofs.
 Section C: Discussion and consistency results on a family of surrogate losses.
 Section D: Bounds on the calibration functions.
 Section D.1: Theorem 7 – a lower bound.
 Section D.2: Theorem 8 – an upper bound.
 Section D.3: Computation of the bounds for specific task losses.
 Section E: Computations of the exact calibration functions for the quadratic surrogate.
 Section E.1: 0-1 loss.
 Section E.2: Block 0-1 loss.
 Section E.3: Hamming loss.
 Section E.4: Mixed 0-1 and block 0-1 loss.
 Section F: Computing constants appearing in the SGD rate.
 Section G: Properties of the basis of the Hamming loss.

A Learning complexity theorem

Theorem 6 (Learning complexity). *Under the assumptions of Theorem 5, for any $\varepsilon > 0$, the random (w.r.t. the observed training set) output $\tilde{f}^{(N)} \in \mathfrak{F}_{F,\mathcal{H}}$ of the ASGD algorithm after*

$$N > N^* := \frac{4D^2M^2}{\check{H}_{\Phi,L,\mathcal{F}}^2(\varepsilon)} \quad (19)$$

iterations has the expected excess risk bounded with ε , i.e., $\mathbb{E}[\mathcal{R}_L(\tilde{f}^{(N)})] < \mathcal{R}_{L,\mathcal{F}}^ + \varepsilon$.*

Proof. By (10) from Theorem 5, N steps of the algorithm, in expectation, result in $\check{H}_{\Phi,L,\mathcal{F}}(\varepsilon)$ accuracy on the surrogate risk, i.e., $\mathbb{E}[\mathcal{R}_\Phi(\tilde{f}^{(N)})] - \mathcal{R}_{\Phi,\mathcal{F}}^* < \check{H}_{\Phi,L,\mathcal{F}}(\varepsilon)$. We now generalize the proof of Theorem 2 to the case of expectation w.r.t. $\tilde{f}^{(N)}$ depending on the random samples used by the ASGD algorithm. We take the expectation of (4) w.r.t. $\tilde{f}^{(N)}$ substituted as \mathbf{f} and use Jensen's inequality (by convexity of $\check{H}_{\Phi,L,\mathcal{F}}$) to get $\mathbb{E}[\mathcal{R}_\Phi(\tilde{f}^{(N)})] - \mathcal{R}_{\Phi,\mathcal{F}}^* \geq \mathbb{E}[\check{H}_{\Phi,L,\mathcal{F}}(\mathcal{R}_L(\tilde{f}^{(N)}) - \mathcal{R}_{L,\mathcal{F}}^*)] \geq \check{H}_{\Phi,L,\mathcal{F}}(\mathbb{E}[\mathcal{R}_L(\tilde{f}^{(N)})] - \mathcal{R}_{L,\mathcal{F}}^*)$. Finally, monotonicity of $\check{H}_{\Phi,L,\mathcal{F}}$ implies $\mathbb{E}[\mathcal{R}_L(\tilde{f}^{(N)})] - \mathcal{R}_{L,\mathcal{F}}^* < \varepsilon$. \square

B Technical lemmas

In this section, we prove two technical lemmas that simplify the proofs of the main theoretical claims of the paper.

Lemma 9 computes the excess of the weighted surrogate risk $\delta\phi$ for the quadratic loss Φ_{quad} (12), which is central to our analysis presented in Section 4. The key property of this result is that the excess $\delta\phi$ is jointly convex w.r.t. the parameters θ and conditional distribution \mathbf{q} , which simplifies further analysis.

Lemma 10 allows to cope with the combinatorial aspect of the computation of the calibration function. In particular, when the excess of the weighted surrogate risk is convex, Lemma 10 reduces the computation of the calibration function to a set of convex optimization problems, which often can be solved analytically. For symmetric losses, such as the 0-1, block 0-1 and Hamming losses, Lemma 10 also provides “symmetry breaking”, meaning that many of the obtained convex optimization problems are identical up to a permutation of labels.

Lemma 9. *Consider the quadratic surrogate Φ_{quad} (12) 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 $F \in \mathbb{R}^{k \times r}$, and assume that $\text{span}(L) \subseteq \mathcal{F}$. Then, the excess of the weighted surrogate loss can be expressed as*

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

Proof. By using the definition of the quadratic surrogate Φ_{quad} (12), we have

$$\begin{aligned} \phi(\mathbf{f}(\theta), \mathbf{q}) &= \frac{1}{2k} (\theta^\top F^\top F \theta + 2\theta^\top F^\top L\mathbf{q}) + r(\mathbf{q}), \\ \theta^* &:= \text{argmin}_{\theta} \phi(\mathbf{f}(\theta), \mathbf{q}) = -(F^\top F)^\dagger F^\top L\mathbf{q}, \\ \delta\phi(\mathbf{f}(\theta), \mathbf{q}) &= \frac{1}{2k} (\theta^\top F^\top F \theta + 2\theta^\top F^\top L\mathbf{q} \\ &\quad + \mathbf{q}^\top L^\top F (F^\top F)^\dagger F^\top L\mathbf{q}), \end{aligned}$$

where $r(\mathbf{q})$ denotes the quantity independent of parameters θ . Note that $P_F := F(F^\top F)^\dagger F^\top$ is the orthogonal projection on the subspace $\text{span}(F)$, so if $\text{span}(L) \subseteq \text{span}(F)$ we have $P_F L = L$, which finishes the proof. \square

Lemma 10. *In the case of a finite number k of labels, 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 written as*

$$H_{\Phi, L, \mathcal{F}}(\varepsilon) = \min_{\substack{i, j \in \text{pred}(\mathcal{F}) \\ i \neq j}} H_{ij}(\varepsilon), \quad (20)$$

where the set $\text{pred}(\mathcal{F}) \subseteq \mathcal{Y}$ is defined as the set of labels that the predictor can predict for some feasible scores and H_{ij} is defined via minimization of the same objective as (5), but w.r.t. a smaller domain:

$$\begin{aligned} H_{ij}(\varepsilon) &= \inf_{\mathbf{f}, \mathbf{q}} \delta\phi(\mathbf{f}, \mathbf{q}), \\ \text{s.t. } &\ell_i(\mathbf{q}) \leq \ell_j(\mathbf{q}) - \varepsilon, \\ &\ell_i(\mathbf{q}) \leq \ell_c(\mathbf{q}), \quad \forall c \in \text{pred}(\mathcal{F}), \\ &f_j \geq f_c, \quad \forall c \in \text{pred}(\mathcal{F}), \\ &\mathbf{f} \in \mathcal{F}, \\ &\mathbf{q} \in \Delta_k. \end{aligned} \quad (21)$$

Here $\ell_c(\mathbf{q}) := (L\mathbf{q})_c$ is the expected loss if predicting label c . Index i represents a label with the smallest expected loss while 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 a 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 , $j \in \text{pred}(\mathcal{F})$, equals the whole set \mathcal{F} . It is possible that sets \mathcal{F}_j do not fully contain their boundary because of the usage of a particular tie-breaking strategy, but their closure can be expressed as $\overline{\mathcal{F}}_j := \{\mathbf{f} \in \mathcal{F} \mid f_j \geq f_c, \forall c \in \text{pred}(\mathcal{F})\}$.

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. $\delta\ell(\mathbf{f}, \mathbf{q}) = \ell_j(\mathbf{q}) - \ell_i(\mathbf{q}) \geq \varepsilon$) is

$$\Delta_{k, i, j, \varepsilon} := \{\mathbf{q} \in \Delta_k \mid \ell_i(\mathbf{q}) \leq \ell_c(\mathbf{q}), \forall c \in \text{pred}(\mathcal{F}); \ell_j(\mathbf{q}) - \ell_i(\mathbf{q}) \geq \varepsilon\},$$

because $\inf_{\mathbf{f}' \in \mathcal{F}} \ell(\mathbf{f}', \mathbf{q}) = \min_{c \in \text{pred}(\mathcal{F})} \ell_c(\mathbf{q})$.

The union of the sets $\{\mathcal{F}_j \times \Delta_{k, i, j, \varepsilon}\}_{i, j \in \text{pred}(\mathcal{F})}$ thus exactly equals the feasibility set of the optimization problem (5)-(6) (note that this is not true for the union of the sets $\{\overline{\mathcal{F}}_j \times \Delta_{k, i, j, \varepsilon}\}_{i, j \in \text{pred}(\mathcal{F})}$, which can be strictly larger), thus we can rewrite the definition of the calibration function as follows:

$$H_{\Phi, L, \mathcal{F}}(\varepsilon) = \min_{\substack{i, j \in \text{pred}(\mathcal{F}) \\ i \neq j}} \inf_{\substack{\mathbf{f} \in \overline{\mathcal{F}}_j, \\ \mathbf{q} \in \Delta_{k, i, j, \varepsilon}}} \delta\phi(\mathbf{f}, \mathbf{q}). \quad (22)$$

To finish the proof, we use Lemma 27 of [49] claiming that the function $\delta\phi(\mathbf{f}, \mathbf{q})$ is continuous w.r.t. both \mathbf{q} and \mathbf{f} , which allows us to substitute sets \mathcal{F}_j in (22) with their closures $\overline{\mathcal{F}}_j$ without changing the value of the infimum. \square

C Consistent surrogate losses

An ideal surrogate should not only be consistent, but also allow efficient optimization, by, e.g., being convex and allowing fast computation of stochastic gradients. In this paper, we study a generalization to arbitrary multi-class losses of a surrogate loss class from Zhang [49, Section 4.4.2]⁸ that satisfies these requirements:

$$\Phi_{a,b}(\mathbf{f}, \mathbf{y}) := \frac{1}{k} \sum_{c=1}^k (L(c, \mathbf{y})a(f_c) + b(f_c)), \quad (23)$$

where $a, b : \mathbb{R} \rightarrow \mathbb{R}$ are convex functions. A generic method to minimize this surrogate is to use any version of the SGD algorithm, while computing the stochastic gradient by sampling \mathbf{y} from the data generating distribution and a label c uniformly. In the case of the quadratic surrogate Φ_{quad} , we proposed instead in the main paper to compute the sum over c analytically instead of sampling c .

Extending the argument from Zhang [49], we show that the surrogates of the form (23) are consistent w.r.t. a task loss L under some sufficient assumptions formalized in Theorem 11.

Theorem 11 (Sufficient conditions for consistency). *The surrogate loss $\Phi_{a,b}$ is consistent w.r.t. a task loss L , i.e., $H_{\Phi_{a,b}, L, \mathbb{R}^k}(\varepsilon) > 0$ for any $\varepsilon > 0$, under the following conditions on the functions $a(f)$ and $b(f)$:*

1. *The functions a and b are convex and differentiable.*
2. *The function $ca(f) + b(f)$ is bounded from below and has a unique global minimizer (finite or infinite) for all $c \in [0, L_{\max}]$.*
3. *The functions $a(f)$ and $\frac{b'(f)}{a'(f)}$ are strictly increasing.*

Proof. Consider an arbitrary conditional probability vector $\mathbf{q} \in \Delta_k$. Assumption 2 then implies that the global minimizer \mathbf{f}^* of the conditional surrogate risk $\phi(\mathbf{f}, \mathbf{q})$ w.r.t. \mathbf{f} is unique. Assumption 1 allows us to set the derivatives to zero and obtain $\frac{b'(f_c^*)}{a'(f_c^*)} = -\ell_c(\mathbf{q})$ where $\ell_c(\mathbf{q}) := (L\mathbf{q})_c$. Assumption 3 then implies that $f_j^* \geq f_i^*$ holds if and only if $\ell_j(\mathbf{q}) \leq \ell_i(\mathbf{q})$.

Now, we will prove by contradiction that $H(\varepsilon) := H_{\Phi_{a,b}, L, \mathbb{R}^k}(\varepsilon) > 0$ for any $\varepsilon > 0$. Assume that for some $\varepsilon > 0$ we have $H(\varepsilon) = 0$. Lemma 10 then implies that for some $i, j \in \mathcal{Y}$, $i \neq j$, we have $H_{ij}(\varepsilon) = 0$. Note that the domain of (21) defining H_{ij} is separable w.r.t. \mathbf{q} and \mathbf{f} . We can now rewrite (21) as

$$H_{ij}(\varepsilon) = \inf_{\mathbf{q} \in \Delta_{k,i,j,\varepsilon}} \delta\phi^*(\mathbf{q}), \text{ where } \delta\phi^*(\mathbf{q}) := \inf_{\mathbf{f} \in \overline{\mathcal{F}}_j} \delta\phi(\mathbf{f}, \mathbf{q}),$$

where $\Delta_{k,i,j,\varepsilon}$ and $\overline{\mathcal{F}}_j$ are defined in the proof of Lemma 10. Lemma 27 of [49] implies that the function $\delta\phi^*(\mathbf{q})$ is a continuous function of \mathbf{q} . Given that $\Delta_{k,i,j,\varepsilon}$ is a compact set, the infimum is achieved at some point $\mathbf{q}^* \in \Delta_{k,i,j,\varepsilon}$. For this \mathbf{q}^* , the global minimum w.r.t. \mathbf{f} exists (Assumption 2). The uniqueness of the global minimum implies that we have $f_j^* = \max_{c \in \mathcal{Y}} f_c^*$. The argument at the beginning of this proof then implies $\ell_j(\mathbf{q}^*) \leq \ell_i(\mathbf{q}^*)$ which contradicts the inequality $\ell_i(\mathbf{q}^*) \leq \ell_j(\mathbf{q}^*) - \varepsilon$ in the definition of $\Delta_{k,i,j,\varepsilon}$. \square

Note that Theorem 11 actually proves that the surrogate $\Phi_{a,b}$ is order-preserving [49], which is a stronger property than consistency.

Below, we give several examples of possible functions $a(f)$, $b(f)$ that satisfy the conditions in Theorem 11 and their corresponding $f^*(\ell) := \operatorname{argmin}_{\mathbf{f} \in \mathbb{R}^k} \phi(\mathbf{f}, \mathbf{q})$ when $\ell := L\mathbf{q}$:

1. If $a(f) = f$, $b(f) = \frac{f^2}{2}$ then $f^*(\ell) = -\ell$, leading to our quadratic surrogate (12).
2. If $a(f) = \frac{1}{L_{\max}}(\exp(f) - \exp(-f))$, $b(f) = \exp(-f)$ then $f^*(\ell) = \frac{1}{2} \log(1 - \frac{1}{L_{\max}}\ell) - \frac{1}{2} \log(\frac{1}{L_{\max}}\ell)$.

⁸Zhang [49] refers to this surrogate as “decoupled unconstrained background discriminative surrogate”. Note the $1/k$ scaling to make $\Phi_{a,b}$ of order 1.

3. If $a(f) = \frac{1}{L_{\max}} f$, $b(f) = \log(1 + \exp(-f))$ then $f^*(\ell) = \log(1 - \frac{1}{L_{\max}} \ell) - \log(\frac{1}{L_{\max}} \ell)$.

In the case of binary classification, these surrogates reduce to L_2 -, exponential, and logistic losses, respectively.

D Bounds on the calibration function

D.1 Lower bound

Theorem 7 (Lower bound on $H_{\Phi_{\text{quad}}}$). *For any task loss L , its quadratic surrogate Φ_{quad} , and a score subspace \mathcal{F} containing the column space of L , the calibration function can be lower bounded:*

$$H_{\Phi_{\text{quad}}, L, \mathcal{F}}(\varepsilon) \geq \frac{\varepsilon^2}{2k \max_{i \neq j} \|P_{\mathcal{F}} \Delta_{ij}\|_2^2} \geq \frac{\varepsilon^2}{4k},$$

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 .

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., $\mathbf{f}(\boldsymbol{\theta}) = F\boldsymbol{\theta}$. Lemma 9 gives us expression (13) for $\delta\phi_{\text{quad}}(F\boldsymbol{\theta}, \mathbf{q})$, which is jointly convex w.r.t. a conditional probability vector \mathbf{q} and parameters $\boldsymbol{\theta}$.

The optimization problem (5)-(6) is non-convex because the constraint (6) 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 10 in App. B allows to bound the calibration function with the minimization w.r.t. selected labels i and j , $H_{\Phi_{\text{quad}}, L, \mathcal{F}}(\varepsilon) \geq \min_{i \neq j} H_{ij}(\varepsilon)$,⁹ where $H_{ij}(\varepsilon)$ is defined as follows:

$$\begin{aligned} H_{ij}(\varepsilon) = \min_{\boldsymbol{\theta}, \mathbf{q}} \quad & \frac{1}{2k} \|F\boldsymbol{\theta} + L\mathbf{q}\|_2^2, \\ \text{s.t.} \quad & (L\mathbf{q})_i \leq (L\mathbf{q})_j - \varepsilon, \\ & (L\mathbf{q})_i \leq (L\mathbf{q})_c, \quad \forall c \in \text{pred}(\mathcal{F}) \\ & (F\boldsymbol{\theta})_j \geq (F\boldsymbol{\theta})_c, \quad \forall c \in \text{pred}(\mathcal{F}) \\ & \mathbf{q} \in \Delta_k. \end{aligned} \tag{24}$$

To obtain a lower bound, we relax (24) by removing some of the constraints and arrive at

$$H_{ij}(\varepsilon) \geq \min_{\boldsymbol{\theta}, \mathbf{q}} \frac{1}{2k} \|F\boldsymbol{\theta} + L\mathbf{q}\|_2^2, \tag{25}$$

$$\text{s.t.} \quad \Delta_{ij}^\top L\mathbf{q} \leq -\varepsilon, \tag{26}$$

$$\Delta_{ij}^\top F\boldsymbol{\theta} \leq 0, \tag{27}$$

where $\Delta_{ij}^\top L\mathbf{q} = (L\mathbf{q})_i - (L\mathbf{q})_j$, $\Delta_{ij}^\top F\boldsymbol{\theta} = (F\boldsymbol{\theta})_i - (F\boldsymbol{\theta})_j$, and $\Delta_{ij} = \mathbf{e}_i - \mathbf{e}_j \in \mathbb{R}^k$ with $\mathbf{e}_c \in \mathbb{R}^k$ being a vector of all zeros with 1 at position c .

The constraint (26) can be readily substituted with equality

$$\Delta_{ij}^\top L\mathbf{q} = -\varepsilon, \tag{28}$$

without changing the minimum because multiplication of both \mathbf{q} and $\boldsymbol{\theta}$ by the constant $\frac{-\varepsilon}{\Delta_{ij}^\top L\mathbf{q}} \in (0, 1]$ preserves feasibility and can only decrease the objective (25).

We now explicitly solve the resulting constraint optimization problem via the KKT optimality conditions. The stationarity constraints give us

$$\frac{1}{k} F^\top (F\boldsymbol{\theta} + L\mathbf{q}) + \mu F^\top \Delta_{ij} = 0, \tag{29}$$

$$\frac{1}{k} L^\top (F\boldsymbol{\theta} + L\mathbf{q}) + \nu L^\top \Delta_{ij} = 0; \tag{30}$$

⁹To simplify the statement of Theorem 7, we removed the constraints $i, j \in \text{pred}(\mathcal{F})$ from Lemma 10 which said that we should consider only the labels that can be predicted with some feasible scores. A potentially tighter lower bound can be obtained by keeping the $i, j \in \text{pred}(\mathcal{F})$ constraint.

the complementary slackness gives $\mu \Delta_{ij}^\top F \boldsymbol{\theta} = 0$ and the feasibility constraints give (28), (27), and $\mu \geq 0$.

Equation (29) allows to compute

$$\boldsymbol{\theta} = -(F^\top F)^\dagger (k\mu F^\top \Delta_{ij} + F^\top L \mathbf{q}). \quad (31)$$

By substituting (31) into (30) and by using the identity (because $L \in \text{span}(F)$):

$$P_F L = F(F^\top F)^\dagger F^\top L = L, \quad (32)$$

we get $(\mu - \nu) L^\top \Delta_{ij} = 0$. If $L^\top \Delta_{ij} = 0$, the problem (25), (27), (28) is infeasible for $\varepsilon > 0$ implying $H_{ij}(\varepsilon) = +\infty$. Otherwise, we have $\mu = \nu$.

From (31) and (32), we also have that:

$$F \boldsymbol{\theta} + L \mathbf{q} = -k\mu P_F \Delta_{ij}. \quad (33)$$

By plugging (31) into the complementary slackness condition and combining with (28), we get

$$\mu^2 k \|P_F \Delta_{ij}\|_2^2 = \mu \varepsilon$$

implying that either $\mu = 0$ or $\mu k \|P_F \Delta_{ij}\|_2^2 = \varepsilon$. In the first case, Eq. (33) implies $F \boldsymbol{\theta} = -L \mathbf{q}$ making satisfying both (28) and (27) impossible. Thus, the later is satisfied implying that the objective (25) is equal to¹⁰

$$\frac{1}{2k} \|F \boldsymbol{\theta} + L \mathbf{q}\|_2^2 = \frac{\varepsilon^2}{2k \|P_F \Delta_{ij}\|_2^2}.$$

Finally, orthogonal projections contract the L_2 -norm, thus $\|P_F \Delta_{ij}\|_2^2 \leq 2$, which gives the second lower bound in the statement of the theorem and finishes the proof. \square

D.2 Upper bound

Theorem 8 (Upper bound on $H_{\Phi_{\text{quad}}}$). *If a loss matrix L with $L_{\max} > 0$ defines a pseudometric⁷ on labels and there are no constraints on the scores, i.e., $\mathcal{F} = \mathbb{R}^k$, then the calibration function for the quadratic surrogate Φ_{quad} can be upper bounded:*

$$H_{\Phi_{\text{quad}}, L, \mathcal{F}}(\varepsilon) \leq \frac{\varepsilon^2}{2k}, \quad 0 \leq \varepsilon \leq L_{\max}.$$

Proof. After applying Lemmas 9 and 10, we arrive at

$$\begin{aligned} H_{ij}(\varepsilon) &= \inf_{\mathbf{f}, \mathbf{q}} \frac{1}{2k} \|\mathbf{f} + L \mathbf{q}\|^2, \\ \text{s.t. } \ell_i(\mathbf{q}) &\leq \ell_j(\mathbf{q}) - \varepsilon, \\ \ell_i(\mathbf{q}) &\leq \ell_c(\mathbf{q}), \quad \forall c \in \mathcal{Y}, \\ f_j &\geq f_c, \quad \forall c \in \mathcal{Y}, \\ \mathbf{f} &\in \mathbb{R}^k, \quad \mathbf{q} \in \Delta_k. \end{aligned} \quad (34)$$

We now consider labels i and j such that $L_{ij} = L_{\max} > 0$ and the point $q_i = \frac{1}{2} + \frac{\varepsilon}{2L_{ij}}$, $q_j = \frac{1}{2} - \frac{\varepsilon}{2L_{ij}}$ (non-negative for $\varepsilon \leq L_{\max}$). We let $q_c = 0$ for $c \notin \{i, j\}$, $f_j = f_i = -\ell_i(\mathbf{q})$ and $f_c = -\ell_c(\mathbf{q})$ for $c \notin \{i, j\}$. We now show that this assignment is feasible.

We have $\ell_j(\mathbf{q}) = q_i L_{ji} + q_j L_{jj} = q_i L_{ji} = q_i L_{ij}$ by symmetry of L . Similarly, $\ell_i(\mathbf{q}) = q_j L_{ij}$ and thus

$$\ell_j(\mathbf{q}) - \ell_i(\mathbf{q}) = L_{ij} \frac{\varepsilon}{L_{ij}} = \varepsilon.$$

We also have

$$\ell_c(\mathbf{q}) - \ell_i(\mathbf{q}) = q_i L_{ci} + q_j L_{cj} - q_j L_{ij} \geq q_j (L_{ic} + L_{cj} - L_{ij}) \geq 0.$$

The first inequality uses $q_i \geq q_j$ and the second inequality uses the fact that L satisfies the triangle inequality (as a pseudometric). Finally, $f_j - f_c = -\ell_i(\mathbf{q}) + \ell_c(\mathbf{q}) \geq 0$.

We thus have shown that the defined point is feasible, so we compute its objective value. We have

$$\frac{1}{2k} \|\mathbf{f} + L \mathbf{q}\|^2 = \frac{1}{2k} (\ell_j(\mathbf{q}) - \ell_i(\mathbf{q}))^2 = \frac{\varepsilon^2}{2k},$$

which completes the proof. \square

¹⁰The possibility $P_F \Delta_{ij} = 0$ is also covered by this equation with the convention that $1/0 = \infty$ (in this case, $\mu^* = \infty$).

D.3 Computation of the lower bounds for specific task losses

0-1 loss. Let L_{01} denote the loss matrix of the 0-1 loss, i.e., $L_{01}(i, j) := [i \neq j]$.² It is convenient to rewrite it with a matrix notation $L_{01} = \mathbf{1}_k \mathbf{1}_k^\top - \mathbf{I}_k$, where $\mathbf{1}_k \in \mathbb{R}^k$ is the vector of all ones and $\mathbf{I}_k \in \mathbb{R}^{k \times k}$ is the identity matrix. We have $\text{rank}(L_{01}) = k$ (for $k \geq 2$), thus $\text{span}(L) = \mathbb{R}^k$. By putting no constraints on the scores, we can easily apply Theorem 7 and obtain the lower bound of $\frac{\varepsilon^2}{4k}$, which is shown to be tight in Proposition 12 of Section E.1.

Block 0-1 loss. We use the symbol $L_{01,b}$ to denote the loss matrix of the block 0-1 loss with b blocks, i.e., $L_{01,b}(i, j) := [i \text{ and } j \text{ are not in the same block}]$. We use s_v to denote the size of block v , $v = 1, \dots, b$, and then $s_1 + \dots + s_b = k$. In the case when all the blocks are of equal sizes, we denote their size by s and have $k = bs$.

With a matrix notation, we have $L_{01,b} = \mathbf{1}_k \mathbf{1}_k^\top - UU^\top$ where the columns of the matrix $U \in \mathbb{R}^{k \times b}$ are indicators of the blocks. We have $\text{rank}(L_{01,b}) = b$ and can simply define $\mathcal{F}_{01,b} := \text{span}(F_{01,b})$ with $F_{01,b} := U$. If we assume that all the blocks have equal size, then we have $U^\top U = s\mathbf{I}_b$ and $\|P_{\mathcal{F}_{01,b}} \Delta_{ij}\|_2^2 = \frac{2}{s}$ if labels i and j belong to different blocks, while $P_{\mathcal{F}_{01,b}} \Delta_{ij} = 0$ if i and j belong to the same block. This leads to the lower bound $\frac{\varepsilon^2}{4b}$, which is shown to be tight in Proposition 14 of Section E.2.

Hamming loss. Consider the (normalized) Hamming loss between tuples of T binary variables, where \hat{y}_t and y_t are the t -th variables of a prediction $\hat{\mathbf{y}}$ and a correct label \mathbf{y} , respectively:

$$\begin{aligned} L_{\text{Ham},T}(\hat{\mathbf{y}}, \mathbf{y}) &:= \frac{1}{T} \sum_{t=1}^T [\hat{y}_t \neq y_t] \\ &= \frac{1}{T} \sum_{t=1}^T ([\hat{y}_t = 0][y_t = 1] + [\hat{y}_t = 1][y_t = 0]) \\ &= \frac{1}{T} \sum_{t=1}^T (1 - [\hat{y}_t = 1])[y_t = 1] + [\hat{y}_t = 1][y_t = 0]) \\ &= \frac{1}{T} \sum_{t=1}^T [y_t = 1] + \frac{1}{T} \sum_{t=1}^T ([y_t = 0] - [y_t = 1]) [\hat{y}_t = 1] \\ &= \alpha_0(\mathbf{y}) + \sum_{t=1}^T \alpha_t(\mathbf{y}) [\hat{y}_t = 1], \end{aligned} \tag{35}$$

The vectors $\alpha_t(\cdot)$ depend only on the column index of the loss matrix. The decomposition (35) implies that $\mathcal{F}_{\text{Ham},T} := \text{span}(F_{\text{Ham},T})$ equals to $\text{span}(L_{\text{Ham},T})$ for $F_{\text{Ham},T} := [\frac{1}{2}\mathbf{1}_{2T}, \mathbf{h}^{(1)}, \dots, \mathbf{h}^{(T)}]$, $(\mathbf{h}^{(t)})_{\hat{\mathbf{y}}} := [\hat{y}_t = 1]$, $t = 1, \dots, T$. We also have that $\text{rank}(L_{\text{Ham},T}) = \text{rank}(F_{\text{Ham},T}) = T + 1$.

In Section G, we show that $\max_{i \neq j} \|P_{\mathcal{F}_{\text{Ham},T}} \Delta_{ij}\|_2^2 = \frac{4T}{2^T}$. By plugging this identity into the lower bound (14), we get $H_{\Phi_{\text{quad}}, L_{\text{Ham},T}, \mathcal{F}_{\text{Ham},T}} \geq \frac{\varepsilon^2}{8T}$, which appears to be tight according to Proposition 15 of Section E.3.

Non-tight cases. In the cases of the block 0-1 loss and the mixed 0-1 and block 0-1 loss (Propositions 13 and 16, respectively), we observe gaps between the lower bound (14) and the exact calibration functions, which shows the limitations of the bound. In particular, it cannot detect level- η consistency for $\eta > 0$ (see Definition 3) and does not change when the loss changes, but the score subspace stays the same.

E Exact calibration functions for quadratic surrogate

This section presents our derivations for the exact values of the calibration functions for different losses. While doing these derivations, we have used numerical simulations and symbolic derivations to check for correctness. Our numerical and symbolic tools are available online.¹¹

E.1 0-1 loss

Proposition 12. Let L_{01} be the 0-1 loss, i.e., $L_{01}(i, j) = [i \neq j]$. Then, the calibration function equals the following quadratic function w.r.t. ε :

$$H_{\Phi_{\text{quad}}, L_{01}, \mathbb{R}^k}(\varepsilon) = \frac{\varepsilon^2}{4k}, \quad 0 \leq \varepsilon \leq 1. \tag{36}$$

¹¹https://github.com/aosokin/consistentSurrogates_derivations

Note that in the case of binary classification, the function (36) is equal to the calibration function for the least squares and truncated least squares surrogates [5, 43].

Proof. First, Lemma 9 with $\mathcal{F} = \mathbb{R}^k$ and $F = \mathbf{I}_k$ gives us the expression

$$\delta\phi_{\text{quad}}(F\boldsymbol{\theta}, \mathbf{q}) = \frac{1}{2k} \|\mathbf{f} + L\mathbf{q}\|_2^2, \quad (37)$$

with $\mathbf{f} = \boldsymbol{\theta} \in \mathbb{R}^k$.

We now reduce the optimization problem (5)-(6) to a convex one by using Lemma 10 and by writing $H_{\Phi_{\text{quad}}, L_{01}, \mathbb{R}^k}(\varepsilon) = \min_{i \neq j \in \mathcal{Y}} H_{ij}(\varepsilon)$, which holds because $\text{pred}(\mathbb{R}^k) = \mathcal{Y}$. Because of the symmetries of the 0-1 loss, all the choices of i and j give the same (up to a permutation of labels) optimization problem to compute $H_{ij}(\varepsilon)$. The definition of the 0-1 loss implies $(L\mathbf{q})_c = 1 - q_c$, which simplifies the excess of the expected task loss appearing in (6) to $\delta\ell(\mathbf{f}, \mathbf{q}) = (L\mathbf{q})_j - (L\mathbf{q})_i = q_i - q_j$. After putting all these together, we get

$$\begin{aligned} H_{ij}(\varepsilon) = \min_{\mathbf{f}, \mathbf{q}} \quad & \frac{1}{2k} \sum_{c=1}^k (f_c + 1 - q_c)^2, \\ \text{s.t.} \quad & q_i \geq q_j + \varepsilon, \\ & q_i \geq q_c, \quad c = 1, \dots, k, \\ & f_j \geq f_c, \quad c = 1, \dots, k, \\ & \sum_{c=1}^k q_c = 1, \quad q_c \geq 0. \end{aligned} \quad (38)$$

We claim that there exists an optimal point of (38), $\mathbf{f}^*, \mathbf{q}^*$, such that $q_c^* = 0$, $c \notin \{i, j\}$, $q_i^* = \frac{1}{2} + \frac{\varepsilon}{2}$, $q_j^* = \frac{1}{2} - \frac{\varepsilon}{2}$; $f_c^* = -1$, $c \notin \{i, j\}$, $f_i^* = f_j^* = f^*$. Note that apart from the specific value of f^* , this is the same point used to prove the upper bound of Theorem 8. After proving this, we will minimize the objective w.r.t. remaining scores at this point.¹²

First, if any $q_c^* = \delta > 0$, $c \notin \{i, j\}$, we can safely move this probability mass to q_i and q_j with the operation

$$\begin{aligned} q_c^* &:= q_c^* - \delta = 0, & q_i^* &:= q_i^* + \frac{\delta}{2}, & q_j^* &:= q_j^* + \frac{\delta}{2}, \\ f_c^* &:= f_c^* - \delta, & f_i^* &:= f_i^* + \frac{\delta}{2}, & f_j^* &:= f_j^* + \frac{\delta}{2}, \end{aligned}$$

which keeps all the constraints of (38) feasible and does not change the objective value.

Second, all the scores f_c^* have to belong to the segment $[-1, 0]$ otherwise clipping them will decrease the objective. With this, setting $f_c^* := -1$, $c \notin \{i, j\}$ can only decrease the objective and will not violate the constraints.

We now show that the equality $q_i^* = q_j^* + \varepsilon$ can hold at the optimum. Indeed, if $q_i^* - q_j^* = \delta' > \varepsilon$, the operation

$$\begin{aligned} q_i^* &:= q_i^* - \frac{\delta' - \varepsilon}{2}, & q_j^* &:= q_j^* + \frac{\delta' - \varepsilon}{2}, \\ f_i^* &:= f_i^* - \frac{\delta' - \varepsilon}{2}, & f_j^* &:= f_j^* + \frac{\delta' - \varepsilon}{2}. \end{aligned} \quad (39)$$

keeps the objective the same and maintains the feasibility constraints. So combining with $q_i^* + q_j^* = 1$, we can now conclude that $q_i^* = \frac{1}{2} + \frac{\varepsilon}{2}$, $q_j^* = \frac{1}{2} - \frac{\varepsilon}{2}$ is an optimal point.

We now show that the equality $f_i^* = f_j^*$ can hold at the optimum. First, we know that the values f_i^* and f_j^* belong to the segment $[q_j^* - 1, q_i^* - 1]$, otherwise we can always truncate the values to the borders of the segment and get an improvement of the objective. Finally, since the inequality $f_j^* \geq f_i^*$ must hold, we conclude that $f_i^* = f_j^* := f^*$ so that f_i^* is closest to its target $q_i^* - 1$ to minimize the objective.

¹²Note that just showing the feasibility of the assigned values \mathbf{q}^* and \mathbf{f}^* give us an upper bound on the calibration function. In the case of the 0-1 loss, it appears that this upper bound matches the lower bound provided by Theorem 7, so we do not need to prove optimality explicitly. However, we still give this proof as a simple illustration of the proof technique as its structure will be re-used also for the cases when the bound of Theorem 7 is not tight.

At the optimal point defined above, it remains to find the value f^* delivering the minimum of the objective. We can achieve this by computing

$$H_{ij}(\varepsilon) = \frac{1}{2k} \min_{f \in [-1, 0]} (f + \frac{1}{2} - \frac{\varepsilon}{2})^2 + (f + \frac{1}{2} + \frac{\varepsilon}{2})^2,$$

which implies $f^* = -0.5$ and $H_{\Phi_{\text{quad}}, L_{01}, \mathbb{R}^k}(\varepsilon) = \frac{\varepsilon^2}{4k}$. \square

Remark. We note that the conditional distribution used in the proof above, $q_i = \frac{1}{2} + \frac{\varepsilon}{2}$, $q_j = \frac{1}{2} - \frac{\varepsilon}{2}$, $q_c = 0$, $c \notin \{i, j\}$, is somewhat unsatisfying from the perspective of explaining why learning the 0-1 loss might be difficult. Indeed, it looks like a gradient based learning algorithm that would start with all values $f_c = -1$ would at the end only optimize over f_i and f_j as the gradient with respect to f_c for $c \notin \{i, j\}$ would stay at zero in $\Phi_{\text{quad}}(\mathbf{f}, \mathbf{y})$ (12) given that only i or j could appear in \mathbf{y} . From this observation, one could think that the calibration function perspective is misleading as SGD could have faster convergence rate than predicted by the worst case for this situation. Fortunately, one can easily check that the point $q_i = \frac{1}{3} + \frac{\varepsilon}{2}$, $q_j = \frac{1}{3} - \frac{\varepsilon}{2}$, $q_c = \frac{1}{3(k-2)}$ for $c \notin \{i, j\}$, $f_i = f_j = \frac{1}{3}$ and $f_c = -\ell_c(\mathbf{q})$ for $c \notin \{i, j\}$ is feasible for (38) and yields the same optimal value of $\frac{\varepsilon^2}{4k}$ for the objective, thus providing another example where the exponential multiclass nature is more readily apparent and cannot be fixed by some “natural initialization” of the learning algorithm.

E.2 Block 0-1 loss

Recall that $L_{01,b}$ is the block 0-1 loss, i.e., $L_{01,b}(i, j) = [i \text{ and } j \text{ are not in the same block}]$. We use b to denote the total number of blocks and s_v to denote the size of block v , $v = 1, \dots, b$. In this section, we compute the calibration functions for the case of unconstrained scores (Proposition 13) and for the case of the scores belonging to the column span of the loss matrix (Proposition 14).

Proposition 13. *Without constraints on the scores, the calibration function for the block 0-1 loss equals the following quadratic function w.r.t. ε :*

$$H_{\Phi_{\text{quad}}, L_{01,b}, \mathbb{R}^k}(\varepsilon) = \frac{\varepsilon^2}{4k} \min_{v=1, \dots, b} \frac{2s_v}{s_v+1} \leq \frac{\varepsilon^2}{2k}, \quad 0 \leq \varepsilon \leq 1.$$

Note that when $s_v = 1$ for some v , we have $H_{\Phi_{\text{quad}}, L_{01,b}, \mathbb{R}^k}(\varepsilon)$ matching to the $\frac{\varepsilon^2}{4k}$ lower bound of Theorem 7. When $s_v \rightarrow \infty$ for all blocks, we have $H_{\Phi_{\text{quad}}, L_{01,b}, \mathbb{R}^k}(\varepsilon)$ matching to the $\frac{\varepsilon^2}{2k}$ upper bound of Theorem 8.

Proof. This proof is of the same structure as the proof of Proposition 12 above.

We use $b(i) \in 1, \dots, b$ to denote the block to which label i belongs and \mathcal{Y}_v to denote the set of labels that belong to block v . We also use Q_v , $v \in 1, \dots, b$, as a shortcut to $\sum_{i \in \mathcal{Y}_v} q_i$, which is the total probability mass on block v .

We start by noting that the i -th component of the vector $(L_{01,b})\mathbf{q}$ equals $1 - Q_{b(i)}$. By applying Lemmas 9, 10, we get

$$H_{ij}(\varepsilon) = \min_{\mathbf{f}, \mathbf{q}} \frac{1}{2k} \sum_{v=1}^b \sum_{c \in \mathcal{Y}_v} (f_c + 1 - Q_{b(c)})^2, \quad (40)$$

$$Q_{b(i)} - Q_{b(j)} \geq \varepsilon, \quad (41)$$

$$Q_{b(i)} \geq Q_u, \quad u = 1, \dots, b,$$

$$f_j \geq f_c, \quad c = 1, \dots, k,$$

$$\sum_{c=1}^k q_c = 1, \quad q_c \geq 0.$$

Analogously to Proposition 12, we claim that there exists an optimal point of (40) such that $q_c = 0$, $c \notin \{i, j\}$; $q_i = 0.5 + \frac{\varepsilon}{2} = Q_{b(i)}$; $q_j = 0.5 - \frac{\varepsilon}{2} = Q_{b(j)}$; $f_c = -1$, $c \notin \mathcal{Y}_{ij} := \mathcal{Y}_{b(i)} \cup \mathcal{Y}_{b(j)}$.

At first, note that if $b(i) = b(j)$, then the constraint (41) is never feasible, so we'll assume that $b(i) \neq b(j)$.

We will now show that we can consider only configurations with all the probability mass on the two selected blocks. Consider some optimal point $\mathbf{f}^*, \mathbf{q}^*$ and denote with $\delta = \sum_{c \in \mathcal{Y} \setminus \mathcal{Y}_{ij}} q_c^*$ the probability mass on the unselected blocks. The operation

$$\begin{aligned} f_c^* &:= f_c^* + \frac{\delta}{2}, \quad c \in \mathcal{Y}_{ij}, & f_c^* &:= -1, \quad c \notin \mathcal{Y}_{ij} \\ q_i^* &:= q_i^* + \frac{\delta}{2}, q_j^* := q_j^* + \frac{\delta}{2}, & q_c^* &:= 0, \quad c \notin \mathcal{Y}_{ij} \end{aligned}$$

can only decrease the objective of (40) because the summands corresponding to the unselected blocks are set to zero. All the constraints stay feasible and the summands corresponding to the selected blocks keep their values.

The probability mass within the block $b(i)$ can be safely moved to q_i^* without changing the objective or violating any constraints. Analogously, the probability mass within the block $b(j)$ can be safely moved to q_j^* . By reusing the operation (39), we can now ensure that $q_i^* = q_j^* + \varepsilon$ and thus that $q_i^* = \frac{1}{2} + \frac{\varepsilon}{2}$ and $q_j^* = \frac{1}{2} - \frac{\varepsilon}{2}$.

At the point defined above, we now minimize the objective (40) w.r.t. $f_c, c \in \mathcal{Y}_{ij}$. At an optimal point, all values $f_c^*, c \in \mathcal{Y}_{ij}$, belong to the segment $[Q_{b(j)}^* - 1, Q_{b(i)}^* - 1]$, otherwise we can always truncate the values to the borders of the segment and get an improvement of the objective. For all the scores $f_c^*, c \neq j$, the following identity holds

$$f_c^* = \begin{cases} Q_{b(c)}^* - 1, & \text{if } Q_{b(c)}^* - 1 < f_j^*, \\ f_j^*. \end{cases} \quad (42)$$

Combining with the segment constraint, it implies that in the block of the label i , we have $f_c^* = f_j^*$, $c \in \mathcal{Y}_{b(i)}$, and, in the block of the label j , we have $f_c^* = Q_{b(j)}^* - 1$, $c \in \mathcal{Y}_{b(j)} \setminus j$.

By plugging the obtained values of q_c^* and f_c^* into (40) and denoting the value $f_j^* + 0.5$ with \tilde{f} , we get

$$\begin{aligned} H_{ij}(\varepsilon) &= \min_{\tilde{f}} \frac{1}{2k} \left(s_{b(i)}(\tilde{f} - \frac{\varepsilon}{2})^2 + (\tilde{f} + \frac{\varepsilon}{2})^2 \right), \\ \text{s.t. } \tilde{f} &\in [-\frac{\varepsilon}{2}, \frac{\varepsilon}{2}]. \end{aligned} \quad (43)$$

By setting the derivative of the objective (43) to zero, we get

$$\tilde{f} = \frac{\varepsilon}{2} \frac{s_{b(i)} - 1}{s_{b(i)} + 1},$$

which belongs to the segment $[-\frac{\varepsilon}{2}, \frac{\varepsilon}{2}]$. We compute the function value at this point:

$$H_{ij}(\varepsilon) = \frac{\varepsilon^2}{4k} \frac{2s_{b(i)}}{s_{b(i)} + 1},$$

which finishes the proof. \square

Proposition 14. *Let the scores \mathbf{f} be piecewise constant on the blocks of the loss, i.e. belong to the subspace $\mathcal{F}_{01,b} = \text{span}(L_{01,b}) \subseteq \mathbb{R}^k$. Then, the calibration function equals the following quadratic function w.r.t. ε :*

$$H_{\Phi_{\text{quad}}, L_{01,b}, \mathcal{F}_{01,b}}(\varepsilon) = \frac{\varepsilon^2}{4k} \min_{v \neq u} \frac{2s_v s_u}{s_v + s_u}, \quad 0 \leq \varepsilon \leq 1.$$

If all the blocks are of the same size, we have $H_{\Phi_{\text{quad}}, L_{01,b}, \mathcal{F}_{01,b}}(\varepsilon) = \frac{\varepsilon^2}{4b}$ where b is the number of blocks.

Proof. The constraints on scores $\mathbf{f} \in \mathcal{F}_{01,b}$ simply imply that the scores within all the blocks are equal. Having this in mind, the proof exactly matches the proof of Proposition 13 until the argument around Eq. (42). Now we cannot set the scores of the block $b(j)$ to different values, and, thus they are all equal to f^* .

By plugging the obtained values of q_c^* and f_c^* into (40) and denoting the value $f_j^* + 0.5$ with \tilde{f} , we get

$$\begin{aligned} H_{ij}(\varepsilon) &= \min_{\tilde{f}} \frac{1}{2k} \left(s_{b(i)}(\tilde{f} - \frac{\varepsilon}{2})^2 + s_{b(j)}(\tilde{f} + \frac{\varepsilon}{2})^2 \right), \\ \text{s.t. } \tilde{f} &\in [-\frac{\varepsilon}{2}, \frac{\varepsilon}{2}]. \end{aligned} \quad (44)$$

By setting the derivative of the objective (44) to zero, we get

$$\tilde{f} = \frac{\varepsilon}{2} \frac{s_{b(i)} - s_{b(j)}}{s_{b(i)} + s_{b(j)}},$$

which belongs to the segment $[-\frac{\varepsilon}{2}, \frac{\varepsilon}{2}]$. We now compute the function value at this point:

$$H_{ij}(\varepsilon) = \frac{\varepsilon^2}{4k} \frac{2s_{b(i)}s_{b(j)}}{s_{b(i)} + s_{b(j)}},$$

which finishes the proof. \square

E.3 Hamming loss

Recall that $L_{\text{Ham}, T}$ is the Hamming loss defined over T binary variables (see Eq. (35) for the precise definition). In this section, we compute the calibration function for the case of the scores belonging to the column span of the loss matrix (Proposition 15).

Proposition 15. *Assume that the scores \mathbf{f} always belong to the column span of the Hamming loss matrix $L_{\text{Ham}, T}$, i.e., $\mathcal{F}_{\text{Ham}, T} = \text{span}(L_{\text{Ham}, T}) \subseteq \mathbb{R}^k$. Then, the calibration function can be computed as follows:*

$$H_{\Phi_{\text{quad}}, L_{\text{Ham}, T}, \mathcal{F}_{\text{Ham}, T}}(\varepsilon) = \frac{\varepsilon^2}{8T}, \quad 0 \leq \varepsilon \leq 1.$$

Proof. We start the proof by applying Lemma 10 and by studying the vector of the expected losses $(L_{\text{Ham}, T})\mathbf{q}$. We note that the $\hat{\mathbf{y}}$ -th element $\ell_{\hat{\mathbf{y}}}(\mathbf{q})$, $\hat{\mathbf{y}} = (\hat{y}_t)_{t=1}^T$, $\hat{y}_t \in \{0, 1\}$, has a simple form of

$$\ell_{\hat{\mathbf{y}}}(\mathbf{q}) = \sum_{\mathbf{y} \in \mathcal{Y}} \frac{q_{\mathbf{y}}}{T} \sum_{t=1}^T [\hat{y}_t \neq y_t] = 1 - \frac{1}{T} \sum_{t=1}^T \sum_{\mathbf{y} \in \mathcal{Y}} q_{\mathbf{y}} [\hat{y}_t = y_t].$$

The quantity $\sum_{\mathbf{y} \in \mathcal{Y}} q_{\mathbf{y}} [\hat{y}_t = y_t]$ corresponds to the marginal probability of a variable t taking a label \hat{y}_t . Note that the expected loss $\ell_{\hat{\mathbf{y}}}(\mathbf{q})$ only depends on \mathbf{q} through marginal probabilities, thus two distributions \mathbf{q}_1 and \mathbf{q}_2 with the same marginals would be indistinguishable when plugged in the optimization problem for $H_{ij}(\varepsilon)$ (21), given that both the constraints and the objective (by Lemma 9) only depend on \mathbf{q} through the expected loss $\ell_{\hat{\mathbf{y}}}(\mathbf{q})$. Having this in mind, we can consider only separable distributions, i.e., $q_{\mathbf{y}} = \prod_{t=1}^T (q_t[y_t = 1] + (1 - q_t)[y_t = 0])$, where $q_t \in [0, 1]$, $t = 1, \dots, T$, are the parameters defining the distribution.

By combining the notation above with Lemmas 9 and 10, we arrive at the following optimization problem:

$$H_{\tilde{\mathbf{y}}\hat{\mathbf{y}}}(\varepsilon) = \min_{\mathbf{f}, \mathbf{q}} \frac{1}{2k} \sum_{\mathbf{y} \in \mathcal{Y}} \left(f_{\mathbf{y}} + 1 - \frac{1}{T} \sum_{t=1}^T q_{t, y_t} \right)^2, \quad (45)$$

$$\text{s.t. } \frac{1}{T} \sum_{t=1}^T (q_{t, \tilde{y}_t} - q_{t, \hat{y}_t}) \geq \varepsilon, \quad (46)$$

$$\frac{1}{T} \sum_{t=1}^T (q_{t, \tilde{y}_t} - q_{t, \hat{y}_t}) \geq 0, \quad \forall \mathbf{y} \in \mathcal{Y}, \quad (47)$$

$$f_{\tilde{\mathbf{y}}} \geq f_{\hat{\mathbf{y}}}, \quad \forall \mathbf{y} \in \mathcal{Y}, \quad (48)$$

$$0 \leq q_t \leq 1, \quad t = 1, \dots, T, \quad (49)$$

$$\mathbf{f} \in \mathcal{F}, \quad (50)$$

where q_{t, y_t} is a shortcut to $q_t[y_t = 1] + (1 - q_t)[y_t = 0]$ and labels $\tilde{\mathbf{y}}$ and $\hat{\mathbf{y}}$ serve as the selected labels i and j , respectively.

The calibration function $H_{\Phi_{\text{quad}}, L_{\text{Ham}, T}, \mathcal{F}_{\text{Ham}, T}}(\varepsilon) = \frac{\varepsilon^2}{8T}$ in the formulation of this proposition matches the lower bound provided by Theorem 7 in Section D.3. Thus, it suffices to construct a feasible w.r.t. (46)-(50) assignment of variables \mathbf{f} , \mathbf{q} and labels $\tilde{\mathbf{y}}$, $\hat{\mathbf{y}}$ such that the objective equals the lower bound.

It suffices to simply set $\tilde{\mathbf{y}}$ to all zeros and $\hat{\mathbf{y}}$ to all ones. In this case, the constraints (46) and (47) take the simplified form:

$$\frac{1}{T} \sum_{t=1}^T (1 - 2q_t) \geq \varepsilon, \quad (51)$$

$$q_t \leq \frac{1}{2}, \quad t = 1, \dots, T. \quad (52)$$

We now set $q_t := \frac{1}{2} - \frac{\varepsilon}{2}$, $t = 1, \dots, T$, and $\mathbf{f} := -\frac{1}{2}\mathbf{1}_k$. This point is clearly feasible when $0 \leq \varepsilon \leq 1$, so it remains to compute the value of the objective. We complete the proof by writing (let w be the count of ones in an assignment \mathbf{y}):

$$\begin{aligned} & \frac{1}{2k} \sum_{\mathbf{y} \in \mathcal{Y}} \left(f_{\mathbf{y}} + 1 - \frac{1}{T} \sum_{t=1}^T q_{t, y_t} \right)^2 = \\ & \frac{1}{2k} \sum_{w=0}^T \binom{T}{w} \left(\frac{1}{2} - \frac{1}{T} \left(w \left(\frac{1}{2} - \frac{\varepsilon}{2} \right) + (T-w) \left(\frac{1}{2} + \frac{\varepsilon}{2} \right) \right) \right)^2 = \\ & \frac{1}{2k} \sum_{w=0}^T \binom{T}{w} \left(\frac{\varepsilon}{2} - \frac{w\varepsilon}{T} \right)^2 = \frac{\varepsilon^2}{2k} \sum_{w=0}^T \binom{T}{w} \left(\frac{1}{4} - \frac{w}{T} + \frac{w^2}{T^2} \right) = \\ & \frac{\varepsilon^2}{2k} \left(\frac{1}{4} 2^T - \frac{1}{T} T 2^{T-1} + \frac{1}{T^2} T(T+1) 2^{T-2} \right) = \frac{\varepsilon^2}{8T}, \end{aligned}$$

where we use the equality $k = 2^T$ and the identities $\sum_{t=0}^T \binom{T}{t} = 2^T$, $\sum_{t=0}^T t \binom{T}{t} = T 2^{T-1}$, $\sum_{t=0}^T t^2 \binom{T}{t} = T(T+1) 2^{T-2}$. \square

E.4 Mixed 0-1 and block 0-1 loss

Recall that $L_{01,b,\eta}$ is the convex combination of the 0-1 loss and the block 0-1 loss with b blocks, i.e., $L_{01,b,\eta} = \eta L_{01} + (1-\eta) L_{01,b}$, $0 \leq \eta \leq 1$. Let all the blocks be of the same size $s = \frac{k}{b} \geq 2$. In this section, we compute the calibration functions for the case of unconstrained scores (Proposition 16) and for the case when scores belong to the column span of the loss matrix (Proposition 17).

Proposition 16. *If there are no constraints on scores \mathbf{f} then the calibration function*

$$H_{\Phi_{\text{quad}}, L_{01,b,\eta}, \mathbb{R}^k}(\varepsilon) = \begin{cases} \frac{\varepsilon^2}{4k}, & \varepsilon \leq \frac{\eta}{1-\eta}, \\ \frac{\varepsilon^2 s}{2k(s+1)} - \frac{\eta(\varepsilon+1)(s-1)}{4k(s+1)} (2\varepsilon - \varepsilon\eta - \eta) & \frac{\eta}{1-\eta} \leq \varepsilon \leq 1 \end{cases}$$

shows that the surrogate is consistent.

Note that when $\eta = 0$, we have $H(\varepsilon) = \frac{\varepsilon^2}{4k} \frac{2s}{s+1}$ as in Proposition 13. When $\eta \geq 0.5$ we have $H(\varepsilon) = \frac{\varepsilon^2}{4k}$, which matches Proposition 12.

Proof. This proof is very similar to the proof of Proposition 13, but technically more involved.

We start by noting that the i -th element of the vector $(L_{01,b,\eta})\mathbf{q}$ equals

$$\sum_{j: b(j) \neq b(i)} (1-\eta)q_j + \sum_{j: j=i} \eta q_j = \eta(1-q_i) + (1-\eta)(1-Q_{b(i)}), \quad (53)$$

where for $b(i)$ and Q_v we reuse the notation defined in the proof of Proposition 13. By combining this with Lemmas 9 and 10, we get

$$\begin{aligned} H_{ij}(\varepsilon) = \min_{\mathbf{f}, \mathbf{q}} & \frac{1}{2k} \sum_{v=1}^b \sum_{c \in \mathcal{Y}_v} (f_c + 1 - \eta q_c - (1-\eta)Q_{b(c)})^2, \\ \text{s.t. } & \eta(q_i - q_j) + (1-\eta)(Q_{b(i)} - Q_{b(j)}) \geq \varepsilon, \\ & \eta(q_i - q_c) + (1-\eta)(Q_{b(i)} - Q_{b(c)}) \geq 0, \forall c \\ & f_j \geq f_c, \forall c, \\ & \sum_{c=1}^k q_c = 1, \quad q_c \geq 0, \forall c. \end{aligned} \quad (54)$$

The blocks are all of the same size so we need to consider just the two cases: 1) the selected labels belong to the same block, i.e., $b(i) = b(j)$; 2) the selected labels belong to the two different blocks, i.e., $b(i) \neq b(j)$.

The first case can be proven by a straight forward generalization of the proof of Proposition 12. Given that the loss value is bounded by 1, the maximal possible value of ε when the constraints can be feasible equals η . Thus, we have $H_{ij}(\varepsilon) = \frac{\varepsilon^2}{4k}$ for $\varepsilon \leq \eta$ and $+\infty$ otherwise.

We will now proceed to the second case $b(i) \neq b(j)$. We show that

$$H_{ij}(\varepsilon) = \begin{cases} \frac{\varepsilon^2}{4k}, & \text{for } \varepsilon \leq \frac{\eta}{1-\eta}, \\ \frac{\varepsilon^2 s}{2k(s+1)} - \frac{\eta(\varepsilon+1)(s-1)}{4k(s+1)}(2\varepsilon - \varepsilon\eta - \eta), & \text{otherwise.} \end{cases}$$

Similarly to the arguments used in Propositions 12 and 13, we claim that there is an optimal point of (54) such that $q_c^* = 0$, $c \notin \{i, j\}$; $q_i^* = 0.5 + \frac{\varepsilon}{2}$; $q_j^* = 0.5 - \frac{\varepsilon}{2}$; and $f_c^* = -1$ for $c \notin \mathcal{Y}_{ij} := \mathcal{Y}_{b(i)} \cup \mathcal{Y}_{b(j)}$.

First, we will show that we can consider only configurations with all the probability mass on the two selected blocks $b(i)$ and $b(j)$. Given any optimal point \mathbf{f}^* and \mathbf{q}^* , the operation (with $\delta = \sum_{c \notin \mathcal{Y}_{ij}} q_c^*$)

$$\begin{aligned} f_i^* &:= f_i^* + \frac{\delta}{2}, & q_i^* &:= q_i^* + \frac{\delta}{2}, \\ f_j^* &:= f_j^* + \frac{\delta}{2}, & q_j^* &:= q_j^* + \frac{\delta}{2}, \\ f_c^* &:= -1, c \notin \mathcal{Y}_{ij} & q_c^* &:= 0, c \notin \mathcal{Y}_{ij} \\ f_c^* &:= f_c^* + (1-\eta)\frac{\delta}{2}, & c &\in \mathcal{Y}_{ij} \setminus \{i, j\} \end{aligned}$$

can only decrease the objective of (54) because the summands corresponding to the unselected $b-2$ blocks are set to zero. All the constraints stay feasible and the values corresponding to the blocks $b(i)$ and $b(j)$ do not change. The last operation is required, because the values $Q_{b(i)}$, $Q_{b(j)}$ change when we change q_i and q_j . Adding $(1-\eta)\frac{\delta}{2}$ to some scores compensates this and cannot violate the constraints because f_j^* goes up by $\frac{\delta}{2} \geq (1-\eta)\frac{\delta}{2}$.

Now we will show that it is possible to move all the mass to the two selected labels i and j . We cannot simply move the mass within one block, but need to create some overflow and move it to another block in a specific way. Consider $\delta := q_a^*$, which is some non-zero mass on a non-selected label of the block $b(i)$. Then, the operation

$$\begin{aligned} f_i^* &:= f_i^* + \delta\frac{\eta}{2}, & q_i^* &:= q_i^* + \delta(1 - \frac{\eta}{2}), \\ f_j^* &:= f_j^* + \delta\frac{\eta}{2}, & q_j^* &:= q_j^* + \delta\frac{\eta}{2}, \\ f_a^* &:= f_a^* + \delta\frac{\eta}{2}(\eta - 3), & q_a^* &:= q_a^* - \delta = 0, \\ f_c^* &:= f_c^* - \delta\frac{\eta}{2}(1 - \eta), & c &\in \mathcal{Y}_i \setminus \{i, a\} \\ f_c^* &:= f_c^* + \delta\frac{\eta}{2}(1 - \eta), & c &\in \mathcal{Y}_j \setminus \{j\} \end{aligned}$$

does not change the objective value of (54) because the quantities $f_c + 1 - \eta q_c - (1-\eta)Q_{b(c)}$, $c \in \mathcal{Y}_{ij}$, stay constant and all the constraints of (54) stay feasible. We repeat this operation for all $a \in \mathcal{Y}_{b(i)} \setminus \{i\}$ and, thus, move all the probability mass within the block $b(i)$ to the label i . In the block $b(j)$, an analogous operation can move all the mass to the label j .

It remains to show that $q_i^* - q_j^* = \varepsilon$. Indeed, if $q_i^* - q_j^* = \delta' > \varepsilon$, the operation analogous to (39)

$$\begin{aligned} f_i^* &:= f_i^* - \frac{\delta' - \varepsilon}{2}, & q_i^* &:= q_i^* - \frac{\delta' - \varepsilon}{2}, \\ f_j^* &:= f_j^* + \frac{\delta' - \varepsilon}{2}, & q_j^* &:= q_j^* + \frac{\delta' - \varepsilon}{2}, \\ f_c^* &:= f_c^* - (1-\eta)\frac{\delta' - \varepsilon}{2}, & c &\in \mathcal{Y}_{b(i)} \setminus \{i\}, \\ f_c^* &:= f_c^* + (1-\eta)\frac{\delta' - \varepsilon}{2}, & c &\in \mathcal{Y}_{b(j)} \setminus \{j\} \end{aligned}$$

can always set $q_i^* - q_j^* = \varepsilon$, and thus $q_i^* = 0.5 + \frac{\varepsilon}{2}$ and $q_j^* = 0.5 - \frac{\varepsilon}{2}$. After this operation, all the scores of the block $b(i)$ go down and all the scores of the block $b(j)$ go up at most as much as f_j^* , so the constraints $f_j \geq f_c$ cannot get violated.

We now proceed with the computation of $H_{ij}(\varepsilon)$. First, we note that convexity and symmetries of (54) implies that all the non-selected scores within each block are equal.¹³ Denote the scores of the

¹³If these optimal scores are not equal, by symmetry, one can obtain the same objective and feasibility by permuting their corresponding values. By taking a uniform convex combination on all permutations, we obtain a point where all the scores are equal, and by convexity, would yield a lower objective value.

non-selected labels of the block $b(i)$ by f'_i , and the scores of the non-selected labels of the block $b(j)$ by f'_j .

Analogous to all the previous propositions, the truncation argument gives us that all the values f_c^* belong to the segment $[-1, -0.5 + \frac{\varepsilon}{2}]$. For all the optimal values f_c^* , $c \neq j$, the following identity holds:

$$f_c^* = \begin{cases} f_j^*, & \text{if } \eta q_c^* + (1 - \eta)Q_{b(c)}^* - 1 \geq f_j^*, \\ \eta q_c^* + (1 - \eta)Q_{b(c)}^* - 1, & \text{otherwise.} \end{cases}$$

Given that f_i^* wants to equal the maximal possible value $-0.5 + \frac{\varepsilon}{2}$, it implies that $f_i^* = f_j^*$. Denote this value by f .

By, plugging the values of q^* and f^* provided above into the objective of (54), we get

$$\begin{aligned} & \frac{1}{2k} \left((f + 0.5 - \frac{\varepsilon}{2})^2 + (s-1)(f'_i + 1 - (1-\eta)(0.5 + \frac{\varepsilon}{2}))^2 + \right. \\ & \left. (f + 0.5 + \frac{\varepsilon}{2})^2 + (s-1)(f'_j + 1 - (1-\eta)(0.5 - \frac{\varepsilon}{2}))^2 \right). \end{aligned} \quad (55)$$

By minimizing (55) without constraints, we get $f^* = -0.5$, $f_i^* = \frac{1}{2}(1 + \varepsilon)(1 - \eta) - 1$, $f_j^* = \frac{1}{2}(1 - \varepsilon)(1 - \eta) - 1$. We now need to compare f_i^* and f_j^* with f^* to satisfy the constraints $f^* \geq f_i^*$ and $f^* \geq f_j^*$. First, we have that

$$f^* - f_j^* = \frac{1}{2}(\eta + \varepsilon - \eta\varepsilon) \geq 0, \text{ for } 0 \leq \varepsilon \leq 1 \text{ and } 0 \leq \eta \leq 1.$$

Second, we have

$$f^* - f_i^* = \frac{1}{2}(\eta - \varepsilon + \eta\varepsilon) \geq 0, \text{ for } 0 \leq \varepsilon \leq \frac{\eta}{1-\eta} \text{ and } 0 \leq \eta \leq 1.$$

We can now conclude that when $\varepsilon \leq \frac{\eta}{1-\eta}$ we have both f'_i and f'_j equal to their unconstrained minimum points leading to $H_{ij}(\varepsilon) = \frac{\varepsilon^2}{4k}$.

Now, consider the case $\varepsilon > \frac{\eta}{1-\eta}$. We have the constraint $f \geq f'_i$ violated, so at the minimum we have $f'_i = f$. The new unconstrained minimum w.r.t. f equals $f^* = \frac{1}{s+1}(-1 - (s-1)(1 - \frac{1}{2}(1-\eta)(1-\varepsilon)))$. We now show that the inequality $f^* \geq f'_j$ still holds. We have

$$f^* - f_j^* = \frac{\eta + \varepsilon s - \eta\varepsilon s}{s+1} \geq 0, \text{ for } 0 \leq \varepsilon \leq 1 \text{ and } 0 \leq \eta \leq 1.$$

Substitution of $f_i^* = f^*$ and f_j^* into (55) gives us

$$\frac{1}{k} \left(\frac{\varepsilon^2 s}{2(s+1)} - \frac{\eta(\varepsilon+1)(s-1)}{4(s+1)} (2\varepsilon - \varepsilon\eta - \eta) \right),$$

which equals $H_{ij}(\varepsilon)$ for $1 \geq \varepsilon > \frac{\eta}{1-\eta}$.

Comparing cases 1 and 2, we observe that $H_{ij}(\varepsilon)$ from case 2 is never larger than the one of case 1, thus case 2 provides the overall calibration function $H_{ij}(\varepsilon)$. \square

Proposition 17. *If the scores f are constrained to be equal inside the blocks, i.e. belong to the subspace $\mathcal{F}_{01,b} = \text{span}(L_{01,b}) \subseteq \mathbb{R}^k$, then the calibration function*

$$H_{\Phi_{\text{quad}}, L_{01,b}, \eta, \mathcal{F}_{01,b}}(\varepsilon) = \begin{cases} \frac{(\varepsilon - \frac{\eta}{2})^2}{4b} \frac{(\frac{\eta}{k} + 1 - \eta)^2}{(1 - \frac{\eta}{2})^2}, & \frac{\eta}{2} \leq \varepsilon \leq 1, \\ 0, & 0 \leq \varepsilon \leq \frac{\eta}{2} \end{cases}$$

shows that the surrogate is consistent up to level $\frac{\eta}{2}$.

When $\eta = 0$, we have $H(\varepsilon) = \frac{\varepsilon^2}{4b}$ as in Proposition 14. When $\eta > 0$ we have $H(\varepsilon) = 0$ for small ε , which corresponds to the case of inconsistent surrogate (0-1 loss and constrained scores).

Proof. This proof combines ideas from Proposition 16 and Proposition 14.

Note that contrary to all the previous results, Lemma 9 is not applicable, because, for $b < k$, we have that $\text{span}(L_{01,b,\eta}) = \mathbb{R}^k \not\subseteq \mathcal{F}_{01,b} = \text{span}(L_{01,b})$.

We now derive an analog of Lemma 9 for this specific case. We define the subspace of scores $\mathcal{F}_{01,b} = \{F\boldsymbol{\theta} \mid \boldsymbol{\theta} \in \mathbb{R}^b\}$ with a matrix $F := F_{01,b} \in \mathbb{R}^{k \times b}$ with columns containing the indicator vectors of the blocks. We have $F^\top F = s\mathbf{I}_b$ and thus $(F^\top F)^{-1} = \frac{1}{s}\mathbf{I}_b$. We shortcut the loss matrix $L_{01,b,\eta}$ to L and rewrite it as

$$L = \eta L_{01} + (1 - \eta)L_{01,b} = \mathbf{1}_k \mathbf{1}_k^\top - \eta \mathbf{I}_k - (1 - \eta)FF^\top.$$

By redoing the derivation of Lemma 9, we arrive at a different excess surrogate:

$$\begin{aligned} \phi(\mathbf{f}(\boldsymbol{\theta}), \mathbf{q}) &= \frac{1}{2k}(s\boldsymbol{\theta}^\top \boldsymbol{\theta} + 2\boldsymbol{\theta}^\top F^\top L\mathbf{q}) + r(\mathbf{q}), \\ \boldsymbol{\theta}^* &:= \operatorname{argmin}_{\boldsymbol{\theta}} \phi(\mathbf{f}(\boldsymbol{\theta}), \mathbf{q}) = -\frac{1}{s}F^\top L\mathbf{q}, \\ \delta\phi(\mathbf{f}(\boldsymbol{\theta}), \mathbf{q}) &= \frac{1}{2k}(s\boldsymbol{\theta}^\top \boldsymbol{\theta} + 2\boldsymbol{\theta}^\top F^\top L\mathbf{q} + \frac{1}{s}\mathbf{q}^\top L^\top FF^\top L\mathbf{q}) \\ &= \frac{s}{2k}\|\boldsymbol{\theta} + \frac{1}{s}F^\top L\mathbf{q}\|_2^2 \\ &= \frac{s}{2k} \sum_{v=1}^s (\theta_v + 1 - (1 - \eta)Q_v - \frac{\eta}{s}Q_v)^2, \end{aligned}$$

where $Q_v = \sum_{c \in \mathcal{Y}_v} q_c$ is the total probability mass on block v and $\mathcal{Y}_v \subset \mathcal{Y}$ denotes the set of labels of block v .

Analogously to Proposition 16 we can now apply Lemma 10 and obtain $H_{ij}(\varepsilon)$.

$$\begin{aligned} H_{ij}(\varepsilon) &= \min_{\boldsymbol{\theta}, \mathbf{q}} \frac{s}{2k} \sum_{v=1}^b (\theta_v + 1 - (1 - \eta)Q_v - \frac{\eta}{s}Q_v)^2, \\ \text{s.t. } &\eta(q_i - q_j) + (1 - \eta)(Q_{b(i)} - Q_{b(j)}) \geq \varepsilon, \\ &\eta(q_i - q_c) + (1 - \eta)(Q_{b(i)} - Q_{b(c)}) \geq 0, \forall c \\ &\theta_{b(j)} \geq \theta_u, \forall u = 1, \dots, b, \\ &\sum_{c=1}^k q_c = 1, \quad q_c \geq 0, \forall c. \end{aligned} \tag{56}$$

The main difference to (54) consists in the fact that we now minimize w.r.t. $\boldsymbol{\theta}$ instead of \mathbf{f} .

Note that because of the way the predictor $\text{pred}(\mathbf{f}(\boldsymbol{\theta}))$ resolves ties (among the labels with maximal scores it always picks the label with the smallest index), not all labels can be predicted. Specifically, only one label from each block can be picked. This argument allows us to assume that $b(i) \neq b(j)$ in the remainder of this proof.

First, let us prove the case for $\varepsilon \leq \frac{\eta}{2}$. We explicitly provide a feasible assignment of variables where the objective equals zero. We set $q_i = \frac{1}{2}$ and $q_c = \frac{1}{2(s-1)}$, $c \in \mathcal{Y}_{b(j)} \setminus \{j\}$. All the other labels (including j and the unselected labels of the block $b(i)$) receive zero probability mass. This assignment of \mathbf{q} implies $Q_{b(i)} = Q_{b(j)} = \frac{1}{2}$ and the zero mass on the other blocks. We also set $\theta_{b(i)}$ and $\theta_{b(j)}$ to $(1 - \eta)\frac{1}{2} + \frac{\eta}{s}\frac{1}{2} - 1$ to ensure zero objective value. Verifying other feasibility constraints we have $\eta(q_i - q_j) + (1 - \eta)(Q_{b(i)} - Q_{b(j)}) = \frac{\eta}{2} \geq \varepsilon$ and $\eta(q_i - q_c) + (1 - \eta)(Q_{b(i)} - Q_{b(c)}) = \eta(\frac{1}{2} - \frac{1}{2(s-1)}) \geq 0$, $c \in \mathcal{Y}_{b(j)} \setminus \{j\}$. Other constraints are trivially satisfied.

Now, consider the case of $\varepsilon > \frac{\eta}{2}$. As usual, we claim the following values of the variables \mathbf{f} and \mathbf{q} result in an optimal point. We have $q_c^* = 0$, $c \notin \mathcal{Y}_{ij}$; $\theta_v^* = -1$, $v \notin \{b(i), b(j)\}$; and $q_i^* = Q_{b(i)}^* = \frac{1+\varepsilon-\eta}{2-\eta}$; $q_c^* = 0$, $c \in \mathcal{Y}_{b(i)} \setminus \{i\}$ (other labels in the block $b(i)$); $q_j^* = 0$, $q_c^* = \frac{1-\varepsilon}{(2-\eta)(s-1)}$, $c \in \mathcal{Y}_{b(j)} \setminus \{j\}$ (other labels in the block $b(j)$).

First, we will show that we can consider only configurations with all the probability mass on the two selected blocks $b(i)$ and $b(j)$. Given some optimal variables \mathbf{f}^* and \mathbf{q}^* , the operation (with $\delta = \sum_{c \in \mathcal{Y} \setminus \mathcal{Y}_{ij}} q_c^*$)

$$\begin{aligned} q_c^* &:= 0, \quad c \in \mathcal{Y} \setminus \mathcal{Y}_{ij}, & q_i^* &:= q_i^* + \frac{\delta}{2}, & q_j^* &:= q_j^* + \frac{\delta}{2}, \\ \theta_v^* &:= -1, \quad v \notin \{b(i), b(j)\}, \\ \theta_{b(i)}^* &:= \theta_{b(i)}^* + \frac{\delta}{2}(1 - \eta + \frac{\eta}{s}), \\ \theta_{b(j)}^* &:= \theta_{b(j)}^* + \frac{\delta}{2}(1 - \eta + \frac{\eta}{s}) \end{aligned}$$

can only decrease the objective of (56) because the summands corresponding to the unselected $b - 2$ blocks are set to zero. All the constraints stay feasible and the values corresponding to the blocks $b(i)$ and $b(j)$ do not change.

Now, we move the mass within the two selected blocks. To start with, moving the mass within one block does not change the objective, because it depends only on $Q_{b(c)}$ and not on q directly. In the block $b(i)$, it is safe to increase q_i and decrease the mass on the other labels, because q_i enters the constraints with the positive sign and while the others enter with the negative sign. So we let $q_c = 0$ for $c \in \mathcal{Y}_{b(i)}/\{i\}$ and $Q_{b(i)} = q_i$. We also have $Q_{b(j)} = 1 - q_i$ as the mass on all other blocks is zero.

Moving mass within the block $b(j)$ is more complicated, as moving mass to some label c of this block might violate the constraints of (56) on q_i . We start by considering the first constraint in (56), using $Q_{b(j)} = 1 - q_i$, we get:

$$q_i \geq \varepsilon + \eta q_j + (1 - \eta)(1 - q_i). \quad (57)$$

By using $q_j \geq 0$ and $\varepsilon \geq \frac{\eta}{2}$, the inequality (57) implies that $q_i \geq \frac{1}{2}$ and thus that

$$q_c \leq Q_{b(j)} \leq \frac{1}{2} \quad \forall c \in \mathcal{Y}_{b(j)}. \quad (58)$$

Now the second constraint of (56) that we want to satisfy is:

$$q_i \geq \eta q_c + (1 - \eta)Q_{b(j)} \quad \forall c \in \mathcal{Y}_{b(j)}. \quad (59)$$

Using (58), we have that the RHS of (59) is $\leq 1/2$, and so since $q_i \geq 1/2$, we have that (59) is satisfied for any valid mass distribution on block $b(j)$ (i.e. such that $Q_{b(j)} \leq 1/2$). Using $q_j = 0$ gives the most possibilities for the value of q_i in the constraint (57). Moreover, the constraint (57) is more stringent than the constraint (59), i.e. if it is satisfied, the second one is also satisfied; so we focus only on the first constraint.

As in the proof of all other propositions, we can make the constraint (57) an equality for the optimum by generalizing the transformation of (39) which makes the constraint tight without changing the objective and maintaining feasibility. So (57) as an equality with $q_j = 0$ yields the value

$$q_i^* = \frac{1+\varepsilon-\eta}{2-\eta}.$$

So to summarize at this point, we have $q_j^* = 0$; $q_c^* = 0$, $c \in \mathcal{Y}_{b(i)} \setminus \{i\}$; $q_c^* = 0$, $\mathcal{Y}_{b(i)} \not\subset \{b(i), b(j)\}$. $q_i^* = \frac{1+\varepsilon-\eta}{2-\eta}$ and $Q_{b(j)} = 1 - q_i^*$. The precise distribution of mass for $c \in \mathcal{Y}_{b(j)}/\{j\}$ does not matter (any distribution is feasible and does not influence the objective, only the total mass matters), but for concreteness, we can choose them to all have the same mass yielding $q_c^* = \frac{1-\varepsilon}{(2-\eta)(s-1)}$, $c \in \mathcal{Y}_{b(j)} \setminus \{j\}$.

We now finish the computation of $H_{ij}(\varepsilon)$. First, we note that, due to the truncation argument similar to the one mentioned in the paragraph after (39), we have both θ_i^* and θ_j^* in the segment $[(1 - \eta)Q_{b(j)}^* + \frac{\eta}{s}Q_{b(j)}^* - 1, (1 - \eta)Q_{b(i)}^* + \frac{\eta}{s}Q_{b(i)}^* - 1]$ and since $\theta_j^* \geq \theta_i^*$, we have $\theta_j^* = \theta_i^* =: \theta$ at the optimum.

Substituting the values $Q_{b(i)}^*$ and $Q_{b(j)}^*$ provided above into the objective of (56) and performing unconstrained minimization w.r.t. θ (we use the help of MATLAB symbolic toolbox to set the derivative to zero) we get

$$\theta^* = -\frac{s-\eta+\eta s}{2s}$$

and, consequently,

$$H_{ij}(\varepsilon) = \frac{s(\varepsilon - \frac{\eta}{2})^2 (\frac{\eta}{s} + 1 - \eta)^2}{4k(1 - \frac{\eta}{2})^2},$$

which finishes the proof. \square

F Constants in the SGD rate

To formalize the learning difficulty by bounding the required number of iterations to get a good value of the risk (Theorem 6), we need to bound the constants D and M . In this section, we provide a

way to bound these constants for the quadratic surrogate Φ_{quad} (12) under a simplifying assumption slightly stronger than the well-specified model Assumption 4.

Consider the family of score functions $\mathfrak{F}_{F,\mathcal{H}}$ defined via an explicit feature map $\psi(\mathbf{x}) \in \mathcal{H}$, i.e., $\mathbf{f}_W(\mathbf{x}) = FW\psi(\mathbf{x})$, where a matrix $F \in \mathbb{R}^{k \times r}$ defines the structure and an operator (which we think of as a matrix with one dimension being infinite) $W : \mathcal{H} \rightarrow \mathbb{R}^r$ contains the learnable parameters. Then the surrogate risk can be written as

$$\mathcal{R}_\Phi(\mathbf{f}_W) = \mathbf{E}_{(\mathbf{x}, \mathbf{y}) \sim \mathcal{D}} \frac{1}{2k} \|FW\psi(\mathbf{x}) + L(:, \mathbf{y})\|_{\mathbb{R}^k}^2$$

and its stochastic w.r.t. (\mathbf{x}, \mathbf{y}) gradient as

$$\mathbf{g}_{\mathbf{x}, \mathbf{y}}(W) = \frac{1}{k} F^\top (FW\psi(\mathbf{x}) + L(:, \mathbf{y})) \psi(\mathbf{x})^\top \quad (60)$$

where $L(:, \mathbf{y})$ denotes the column of the loss matrix corresponding to the correct label \mathbf{y} . Note that computing the stochastic gradient requires performing products $F^\top F$ and $F^\top L(:, \mathbf{y})$ for which direct computation is intractable when k is exponential, but which can be done in closed form for the structured losses we consider (the Hamming and block 0-1 loss). More generally, these operations require suitable inference algorithms.

To derive the constants, we use a simplifying assumption stronger than Assumption 4 in the case of quadratic surrogate: we assume that the conditional $q_c(\mathbf{x})$, seen as a function of \mathbf{x} , belongs to the RKHS \mathcal{H} , which by the reproducing property implies that for each $c = 1, \dots, k$, there exists $v_c \in \mathcal{H}$ such that $q_c(\mathbf{x}) = \langle v_c, \psi(\mathbf{x}) \rangle_{\mathcal{H}}$ for all $\mathbf{x} \in \mathcal{X}$. Concatenating all v_c , we get an operator $V : \mathcal{H} \rightarrow \mathbb{R}^k$. To derive the bound, we also assume that $\sum_{c=1}^k \|v_c\|_{\mathcal{H}} \leq Q_{\max}$ and $\|\psi(\mathbf{x})\|_{\mathcal{H}} \leq R$ for all $\mathbf{x} \in \mathcal{X}$. In the following, we use the notation $\mathbf{q}_{\mathbf{x}}$ to denote the vector in \mathbb{R}^k with components $q_c(\mathbf{x})$, $c = 1, \dots, k$, for a fixed \mathbf{x} , and thus $\mathbf{q}_{\mathbf{x}} = V\psi(\mathbf{x})$.

Under these assumptions, we can write the theoretical minimum of the surrogate risk. The gradient of the surrogate risk gives

$$\begin{aligned} k \nabla_W \mathcal{R}_\Phi(\mathbf{f}_W) &= F^\top F W \mathbf{E}_{\mathbf{x} \sim \mathcal{D}_{\mathcal{X}}} (\psi(\mathbf{x}) \psi(\mathbf{x})^\top) + F^\top L \mathbf{E}_{\mathbf{x} \sim \mathcal{D}_{\mathcal{X}}} (\mathbf{q}_{\mathbf{x}} \psi(\mathbf{x})^\top) \\ &= F^\top F W \mathbf{E}_{\mathbf{x} \sim \mathcal{D}_{\mathcal{X}}} (\psi(\mathbf{x}) \psi(\mathbf{x})^\top) + F^\top L V \mathbf{E}_{\mathbf{x} \sim \mathcal{D}_{\mathcal{X}}} (\psi(\mathbf{x}) \psi(\mathbf{x})^\top) \\ &= (F^\top F W + F^\top L V) \mathbf{E}_{\mathbf{x} \sim \mathcal{D}_{\mathcal{X}}} (\psi(\mathbf{x}) \psi(\mathbf{x})^\top). \end{aligned}$$

Setting the content of the parenthesis to zero gives that $W^* = -(F^\top F)^\dagger F^\top L V$ is a solution to the stationary condition equation $\nabla_W \mathcal{R}_\Phi(\mathbf{f}_W) = 0$.

We can now bound the Hilbert-Schmidt norm of this choice of optimal parameters W^* as

$$\begin{aligned} \|W^*\|_{HS} &= \|(F^\top F)^\dagger F^\top L V\|_{HS} \\ &\leq \|(F^\top F)^\dagger F^\top\|_{HS} \|L V\|_{HS} && // \text{submultiplicativity of } \|\cdot\|_{HS} \\ &\leq \sqrt{r} \|(F^\top F)^\dagger F^\top\|_2 \|L V\|_{HS} && // \text{connection of } \|\cdot\|_{HS} \text{ and } \|\cdot\|_2 \text{ via } r = \text{rank}(F) \\ &= \sqrt{r} \sigma_{\min}^{-1}(F) \|L V\|_{HS} && // \text{rotation invariance of } \|\cdot\|_2 \\ &\leq \sqrt{r} \sigma_{\min}^{-1}(F) \sqrt{k} L_{\max} Q_{\max} =: D && // \text{the definition of } \|\cdot\|_{HS} \text{ and triangular inequality} \end{aligned}$$

where $\|\cdot\|_{HS}$ and $\|\cdot\|_2$ denote the Hilbert-Schmidt and spectral norms, respectively, and $\sigma_{\min}^{-1}(F)$ stands for the smallest singular value of the matrix F . The last inequality follows from the definition of the Hilbert-Schmidt norm $\|L V\|_{HS}^2 = \sum_{i=1}^k \|\sum_{c=1}^k L(i, c) v_c\|_{\mathcal{H}}^2$ and from the triangular inequality $\|\sum_{c=1}^k L(i, c) v_c\|_{\mathcal{H}} \leq \sum_{c=1}^k |L(i, c)| \|v_c\|_{\mathcal{H}} \leq L_{\max} Q_{\max}$ thus giving $\|L V\|_{HS} \leq \sqrt{k} L_{\max} Q_{\max}$.

Analogously, we now bound the Hilbert-Schmidt norm of the stochastic gradient $\mathbf{g}_{\mathbf{x}, \mathbf{y}}(W)$.

$$\begin{aligned} \|\mathbf{g}_{\mathbf{x}, \mathbf{y}}(W)\|_{HS} &\leq \frac{1}{k} \|F^\top F W \psi(\mathbf{x}) + F^\top L(:, \mathbf{y})\|_2 \|\psi(\mathbf{x})\|_{\mathcal{H}} \\ &\leq \frac{1}{k} (\|F^\top F W \psi(\mathbf{x})\|_2 + \|F^\top L(:, \mathbf{y})\|_2) \|\psi(\mathbf{x})\|_{\mathcal{H}} \\ &\leq \frac{1}{k} (\|F^\top F\|_2 \|W\|_{HS} \|\psi(\mathbf{x})\|_{\mathcal{H}} + \|F\|_2 \|L(:, \mathbf{y})\|_2) \|\psi(\mathbf{x})\|_{\mathcal{H}} \\ &\leq \frac{1}{k} \sigma_{\max}^2(F) D R^2 + \frac{1}{k} \sigma_{\max}(F) \sqrt{k} L_{\max} R =: M \end{aligned}$$

where R is an upper bound on $\|\psi(\mathbf{x})\|_{\mathcal{H}}$ and $\sigma_{\max}(F)$ is a maximal singular value of F . Here the first inequality follows from the fact that the rank of $\mathbf{g}_{\mathbf{x}, \mathbf{y}}(W)$ equals 1 and from submultiplicativity of

the spectral norm. We also use the inequality $\|W\psi(\mathbf{x})\|_2 \leq \|W\|_{HS}\|\psi(\mathbf{x})\|_{\mathcal{H}}$, which follows from the properties of the Hilbert-Schmidt norm.

The bound of Theorem 5 contains the quantity DM and the step size of ASGD depends on $\frac{D}{M}$, so, to be practical, both quantities cannot be exponential (for numerical stability; but the important quantity is the number of iterations from Theorem 6). We have

$$DM = \kappa^2(F)R^2rL_{\max}^2Q_{\max}^2 + \kappa(F)R\sqrt{r}L_{\max}^2Q_{\max} = L_{\max}^2\xi(\kappa(F)\sqrt{r}RQ_{\max}), \quad \xi(z) = z^2 + z, \\ \frac{D}{M} = \frac{\sigma_{\max}^2(F)}{k}R^2 + \frac{\sigma_{\max}(F)\sigma_{\min}(F)}{k}\frac{R}{Q_{\max}\sqrt{r}}$$

where $\kappa(F) = \frac{\sigma_{\max}}{\sigma_{\min}}$ is the condition number of F . Note that the quantity DM is invariant to the scaling of the matrix F . The quantity $\frac{D}{M}$ scales proportionally to the square of the scale of F and thus rescaling F can always bring it to $\mathcal{O}(1)$. For the rest of the analysis, we consider R and Q_{\max} to be well-behaved constants and thus focus on the dependence of the quantity DM on F and L .

F.1 Constants for specific losses

We now estimate the product DM from (18) for the 0-1, block 0-1 and Hamming losses. For the definition of the losses and the corresponding matrices F , we refer to Section D.3.

0-1 loss. For the 0-1 loss L_{01} and $F = \mathbf{I}_k$, we have $L_{\max} = 1$, $r = k$, $\sigma_{\min} = \sigma_{\max} = 1$, thus $DM = \mathcal{O}(k)$ is very large leading to very slow convergence of ASGD.

Block 0-1 loss. For the block 0-1 loss $L_{01,b}$ and matrix $F_{01,b}$, we have $L_{\max} = 1$, $r = b$, $\sigma_{\min} = \sigma_{\max} = \sqrt{b}$, thus $DM = \mathcal{O}(b)$.

Hamming loss. For the Hamming loss, we have $L_{\max} = 1$, $r = \log_2 k + 1$, $\kappa(F_{\text{Ham},T}) \leq \log_2 k + 2$ (see the derivation in Section G). Finally, we have $DM = \mathcal{O}(\log_2^3 k)$.

G Properties of the basis of the Hamming loss

As defined in (35), the matrix $L_{\text{Ham},T} \in \mathbb{R}^{k \times k}$ is the matrix of the Hamming loss between tuples of T binary variables, and the number of labels equals $k = 2^T$. Also recall that $F_{\text{Ham},T} := [\frac{1}{2}\mathbf{1}_{2^T}, \mathbf{h}^{(1)}, \dots, \mathbf{h}^{(T)}]$, $(\mathbf{h}^{(t)})_{\hat{\mathbf{y}}} := [\hat{y}_t = 1]$, $t = 1, \dots, T$. We have $\mathcal{F}_{\text{Ham},T} = \text{span}(F_{\text{Ham},T}) = \text{span}(L_{\text{Ham},T})$ and $\text{rank}(L_{\text{Ham},T}) = \text{rank}(F_{\text{Ham},T}) = T + 1$.

We now explicitly compute $\max_{i \neq j} \|P_{\mathcal{F}_{\text{Ham},T}} \Delta_{ij}\|_2^2$. We shortcut $F_{\text{Ham},T}$ by F and compute

$$F^T F = 2^{T-2} \begin{bmatrix} 1 & 1 & \dots & 1 \\ 1 & 2 & 1 & \dots \\ 1 & 1 & 2 & \dots \\ \dots & \dots & \dots & 1 \\ 1 & \dots & 1 & 2 \end{bmatrix}. \quad (61)$$

We can compute the inverse matrix explicitly as well:

$$(F^T F)^{-1} = 2^{2-T} \begin{bmatrix} 1+T & -1 & \dots & -1 \\ -1 & 1 & 0 & \dots \\ -1 & 0 & 1 & \dots \\ \dots & \dots & \dots & 0 \\ -1 & \dots & 0 & 1 \end{bmatrix}. \quad (62)$$

The vector $F^T \Delta_{ij}$ equals the difference of the two rows of F , i.e., $[0, c_1, \dots, c_T]^T \in \mathbb{R}^{T+1}$ with each $c_t \in \{-1, 0, +1\}$. We explicitly compute the square norm $\|P_{\mathcal{F}_{\text{Ham},T}} \Delta_{ij}\|_2^2$:

$$\|P_{\mathcal{F}_{\text{Ham},T}} \Delta_{ij}\|_2^2 = \Delta_{ij}^T F(F^T F)^{-1} F^T \Delta_{ij} = [0, c_1, \dots, c_T](F^T F)^{-1}[0, c_1, \dots, c_T]^T = 2^{2-T} \sum_{t=1}^T c_t^2,$$

where the last equality follows from the identity submatrix of (62) and from the zero in the first position of the vector $F^T \Delta_{ij}$. The quantity $\|P_{\mathcal{F}_{\text{Ham},T}} \Delta_{ij}\|_2^2$ is maximized when none of c_t equals zero, which is achievable, e.g., when the label i corresponds to all zeros and the label j to all ones. We now have $\max_{i \neq j} \|P_{\mathcal{F}_{\text{Ham},T}} \Delta_{ij}\|_2^2 = \frac{4^T}{2^T}$.

We now compute the smallest and largest eigenvalues of the Gram matrix (61) for $F_{\text{Ham},T}$. Ignoring the scaling factor 2^{T-2} , we see by Gaussian elimination that the determinant and thus the product of all eigenvalues equals 1. If we subtract \mathbf{I}_{T+1} the matrix becomes of rank 2, meaning that $T - 1$ eigenvalues equal 1. The trace, i.e., the sum of the eigenvalues of (61), without the scaling factor 2^{T-2} equals $2T + 1$. Summing up, we have $\lambda_{\min}\lambda_{\max} = 1$ and $\lambda_{\min} + \lambda_{\max} = T + 2$. We can now compute $\lambda_{\min} = \frac{1}{2}(T + 2 - \sqrt{T^2 + 4T}) \in [\frac{1}{T+2}, \frac{1}{T}]$ and $\lambda_{\max} = \frac{1}{2}(T + 2 + \sqrt{T^2 + 4T}) \in [T + 1, T + 2]$. By putting back the multiplicative factor, we get $\sigma_{\min} = \sqrt{\lambda_{\min}} \geq \frac{\sqrt{k}}{2\sqrt{\log_2 k + 2}}$ and $\sigma_{\max} = \sqrt{\lambda_{\max}} \leq \frac{\sqrt{k}}{2}\sqrt{\log_2 k + 2}$, and thus the condition number is $\kappa \leq \log_2 k + 2$.