

Learning with Fenchel-Young Losses

Mathieu Blondel

NTT Communication Science Laboratories
Kyoto, Japan

mathieu@mblondel.org

André F. T. Martins

Unbabel & Instituto de Telecomunicações
Lisbon, Portugal

andre.martins@unbabel.com

Vlad Niculae

Instituto de Telecomunicações
Lisbon, Portugal

vlad@vene.ro

Abstract

Over the past decades, numerous loss functions have been proposed for a variety of supervised learning tasks, including regression, classification, ranking, and more generally structured prediction. Understanding the core principles and theoretical properties underpinning these losses is key to choose the right loss for the right problem, as well as to create new losses which combine their strengths. In this paper, we introduce Fenchel-Young losses, a generic way to *construct* a convex loss function for a regularized prediction function. We provide an in-depth study of their properties in a very broad setting, covering all the aforementioned supervised learning tasks, and revealing new connections between sparsity, generalized entropies, and separation margins. We show that Fenchel-Young losses unify many well-known loss functions and allow to create useful new ones easily. Finally, we derive efficient predictive and training algorithms, making Fenchel-Young losses appealing both in theory and practice.

Keywords: loss functions, output regularization, convex duality, structured prediction

1. Introduction

Loss functions are a cornerstone of statistics and machine learning: They measure the difference, or “loss,” between a ground truth and a prediction. As such, much work has been devoted to designing loss functions for a variety of supervised learning tasks, including regression (Huber, 1964), classification (Crammer and Singer, 2001), ranking (Joachims, 2002) and structured prediction (Lafferty et al., 2001; Collins, 2002; Tsochantaridis et al., 2005), to name only a few well-known directions.

For the case of probabilistic classification, proper composite loss functions (Reid and Williamson, 2010; Williamson et al., 2016) offer a principled framework unifying various existing loss functions. A proper composite loss is the composition of a proper loss between two probability distributions, with an invertible mapping from real vectors to probability

distributions. The theoretical properties of proper loss functions, also known as proper scoring rules (Grünwald and Dawid (2004); Gneiting and Raftery (2007); and references therein), such as their Fisher consistency (classification calibration) and correspondence with Bregman divergences, are now well-understood. However, not all existing losses are proper composite loss functions; a notable example is the hinge loss used in support vector machines. In fact, we shall see that any loss function enjoying a separation margin, a prevalent concept in statistical learning theory which has been used to prove the famous perceptron mistake bound (Rosenblatt, 1958) and many other generalization bounds (Vapnik, 1998; Schölkopf and Smola, 2002), cannot be written in composite proper loss form.

At the same time, loss functions are often intimately related to an underlying statistical model and prediction function. For instance, the logistic loss corresponds to the multinomial distribution and the softmax operator, while the conditional random field (CRF) loss (Lafferty et al., 2001) for structured prediction is tied with marginal inference (Wainwright and Jordan, 2008). Both are instances of generalized linear models (Nelder and Baker, 1972; McCullagh and Nelder, 1989), associated with exponential family distributions. More recently, Martins and Astudillo (2016) proposed a new classification loss based on the projection onto the simplex. Unlike the logistic loss, this “sparsemax” loss induces probability distributions with **sparse** support, which is desirable in some applications for interpretability or computational efficiency reasons. However, the sparsemax loss was derived in a relatively ad-hoc manner and it is still relatively poorly understood. Is it one of a kind or can we generalize it in a principled manner? Thorough understanding of the core principles underpinning existing losses and their associated predictive model, potentially enabling the creation of useful new losses, is one of the main quests of this paper.

This paper. The starting point of this paper are the notions of **output** regularization and regularized prediction functions, which we use to provide a variational perspective on many existing prediction functions, including the aforementioned softmax, sparsemax and marginal inference. Based on simple convex duality arguments, we then introduce Fenchel-Young losses, a new way to automatically construct a loss function associated with any regularized prediction function. As we shall see, our proposal recovers many existing loss functions, which is in a sense surprising since many of these losses were originally proposed by independent efforts. Our framework goes beyond the simple probabilistic classification setting: We show how to create loss functions over various structured domains, including convex polytopes and convex cones. Our framework encourages the loss designer to think **geometrically** about the outputs desired for the task at hand. Once a (regularized) prediction function has been designed, our framework generates a corresponding loss function automatically. We will demonstrate the ease of creating loss functions, including useful new ones, using abundant examples throughout this paper.

Organization and contributions. The rest of this paper is organized as follows.

§2 We introduce regularized prediction functions, unifying and generalizing softmax, sparsemax and marginal inference, among many others.

- §3 We introduce Fenchel-Young losses to learn models whose output layer is a regularized prediction function. We provide an in-depth study of their properties, showing that they unify many existing losses, including unstructured and structured losses.
- §4 We study Fenchel-Young losses for probabilistic classification. We show how to seamlessly create entire new families of losses from generalized entropies.
- §5 We characterize which entropies yield sparse distributions and losses with a separation margin, notions we prove to be intimately connected. Furthermore, we show that losses that enjoy a margin and induce a sparse distribution are precisely the ones that cannot be written in proper composite loss form.
- §6 We study Fenchel-Young losses for positive measures (unnormalized probability distributions), providing a new perspective on one-vs-all loss reductions.
- §7 We study Fenchel-Young losses for structured prediction. We introduce the concepts of probability distribution and mean regularizations, providing a unifying perspective on a number of structured losses, including the conditional random field (CRF) loss and the recently-proposed SparseMAP loss (Niculae et al., 2018). We illustrate these results by deriving losses over various convex polytopes, including ranking losses.
- §8 We present primal and dual algorithms for learning models with Fenchel-Young losses defined over arbitrary domains. We derive new efficient algorithms to compute regularized prediction functions and proximity operators, which are a core sub-routine for dual training algorithms.
- §9 We demonstrate the ability of Fenchel-Young losses to induce sparse distributions on two tasks: label proportion estimation and dependency parsing.
- §10 Finally, we review related work on proper (composite) losses and other losses proposed in the literature, Fenchel duality, and approximate inference.

Previous papers. This paper builds upon two previously published shorter conference papers. The first (Niculae et al., 2018) introduced Fenchel-Young losses in the structured prediction setting but only provided a limited analysis of their properties. The second (Blondel et al., 2019) provided a more in-depth analysis but focused on unstructured probabilistic classification. This paper provides a comprehensive study of Fenchel-Young losses across various domains. Besides a much more thorough treatment of previously covered topics, this paper contributes entirely new sections, including §6 on losses for positive measures, §8 on primal and dual training algorithms, and §A.2 on loss “Fenchel-Youngization”. We provide in §7 a new unifying view between structured prediction losses, and discuss at length various convex polytopes, promoting a geometric approach to structured prediction loss design; we also provide novel results in this section regarding structured separation margins (Proposition 8), proving the unit margin of the SparseMAP loss. We demonstrate how to use our framework to create useful new losses, including ranking losses, not covered in the previous two papers.

Notation. We denote the $(d-1)$ -dimensional probability simplex by $\Delta^d := \{\mathbf{p} \in \mathbb{R}_+^d : \|\mathbf{p}\|_1 = 1\}$. We denote the convex hull of a set \mathcal{Y} by $\text{conv}(\mathcal{Y})$ and the conic hull by $\text{cone}(\mathcal{Y})$. We

denote the domain of a function $\Omega: \mathbb{R}^d \rightarrow \mathbb{R} \cup \{\infty\}$ by $\text{dom}(\Omega) := \{\boldsymbol{\mu} \in \mathbb{R}^d: \Omega(\boldsymbol{\mu}) < \infty\}$. We denote the Fenchel conjugate of Ω by $\Omega^*(\boldsymbol{\theta}) := \sup_{\boldsymbol{\mu} \in \text{dom}(\Omega)} \langle \boldsymbol{\theta}, \boldsymbol{\mu} \rangle - \Omega(\boldsymbol{\mu})$. We denote the indicator function of a set \mathcal{C} by

$$I_{\mathcal{C}}(\boldsymbol{\mu}) := \begin{cases} 0 & \text{if } \boldsymbol{\mu} \in \mathcal{C} \\ \infty & \text{otherwise} \end{cases} \quad (1)$$

and its support function by $\sigma_{\mathcal{C}}(\boldsymbol{\theta}) := \sup_{\boldsymbol{\mu} \in \mathcal{C}} \langle \boldsymbol{\theta}, \boldsymbol{\mu} \rangle$. We define the proximity operator (a.k.a. proximal operator) of Ω by

$$\text{prox}_{\Omega}(\boldsymbol{\eta}) := \underset{\boldsymbol{\mu} \in \text{dom}(\Omega)}{\text{argmin}} \frac{1}{2} \|\boldsymbol{\mu} - \boldsymbol{\eta}\|^2 + \Omega(\boldsymbol{\mu}). \quad (2)$$

We denote the interior and relative interior of \mathcal{C} by $\text{int}(\mathcal{C})$ and $\text{relint}(\mathcal{C})$, respectively. We denote $[\boldsymbol{x}]_+ := \max(\boldsymbol{x}, \mathbf{0})$, evaluated element-wise.

2. Regularized prediction functions

In this section, we introduce the concept of regularized prediction function (§2.1), which is central to this paper. We then give simple and well-known examples of such functions (§2.2) and discuss their properties in a general setting (§2.3, §2.4).

2.1 Definition

We consider a general predictive setting with input variables $\boldsymbol{x} \in \mathcal{X}$, and a parametrized model $\boldsymbol{f}_W: \mathcal{X} \rightarrow \mathbb{R}^d$ (which could be a linear model or a neural network), producing a score vector $\boldsymbol{\theta} := \boldsymbol{f}_W(\boldsymbol{x}) \in \mathbb{R}^d$. In a simple multi-class classification setting, the score vector is typically used to pick the highest-scoring class among d possible ones

$$\hat{\boldsymbol{y}}(\boldsymbol{\theta}) \in \underset{j \in [d]}{\text{argmax}} \theta_j. \quad (3)$$

This can be generalized to an arbitrary output space $\mathcal{Y} \subseteq \mathbb{R}^d$ by using instead

$$\hat{\boldsymbol{y}}(\boldsymbol{\theta}) \in \underset{\boldsymbol{y} \in \mathcal{Y}}{\text{argmax}} \langle \boldsymbol{\theta}, \boldsymbol{y} \rangle, \quad (4)$$

where intuitively $\langle \boldsymbol{\theta}, \boldsymbol{y} \rangle$ captures the affinity between \boldsymbol{x} (since $\boldsymbol{\theta}$ is produced by $\boldsymbol{f}_W(\boldsymbol{x})$) and \boldsymbol{y} . Therefore, (4) seeks the output \boldsymbol{y} with greatest affinity with \boldsymbol{x} . The support function $\sigma_{\mathcal{Y}}(\boldsymbol{\theta}) = \max_{\boldsymbol{y} \in \mathcal{Y}} \langle \boldsymbol{\theta}, \boldsymbol{y} \rangle = \langle \boldsymbol{\theta}, \hat{\boldsymbol{y}}(\boldsymbol{\theta}) \rangle$ can be interpreted as the largest projection of any element of \mathcal{Y} onto the line generated by $\boldsymbol{\theta}$.

Clearly, (4) recovers (3) with $\mathcal{Y} = \{\boldsymbol{e}_1, \dots, \boldsymbol{e}_d\}$, where \boldsymbol{e}_j is a standard basis vector, $\boldsymbol{e}_j := [0, \dots, 0, \underbrace{1}_j, 0, \dots, 0]$. In this case, the cardinality $|\mathcal{Y}|$ and the dimensionality d coincide,

but this need not be the case in general. Eq. (4) is often called a linear maximization oracle or maximum a-posteriori (MAP) oracle (Wainwright and Jordan, 2008). The latter name comes from the fact that (4) coincides with the mode of the Gibbs distribution defined by

$$p(\boldsymbol{y}; \boldsymbol{\theta}) \propto \exp \langle \boldsymbol{\theta}, \boldsymbol{y} \rangle.$$

Prediction over convex hulls. We now extend the prediction function (4) by replacing \mathcal{Y} with its convex hull $\text{conv}(\mathcal{Y}) := \{\mathbb{E}_{\mathbf{p}}[Y] : \mathbf{p} \in \Delta^{|\mathcal{Y}|}\}$ and introducing a regularization function Ω into the optimization problem:

$$\hat{\mathbf{y}}_{\Omega}(\boldsymbol{\theta}) \in \underset{\boldsymbol{\mu} \in \text{conv}(\mathcal{Y})}{\text{argmax}} \langle \boldsymbol{\theta}, \boldsymbol{\mu} \rangle - \Omega(\boldsymbol{\mu}).$$

We emphasize that the regularization is w.r.t. predictions (outputs) and not w.r.t. model parameters (denoted by W in this paper), as is usually the case in the literature. We illustrate the regularized prediction function pipeline in Figure 1.

Unsurprisingly, the choice $\Omega = 0$ recovers the unregularized prediction function (4). This follows from the fundamental theorem of linear programming (Dantzig et al., 1955, Theorem 6), which states that the maximum of a linear form over a convex polytope is always achieved at one of its vertices:

$$\hat{\mathbf{y}}_0(\boldsymbol{\theta}) \in \underset{\boldsymbol{\mu} \in \text{conv}(\mathcal{Y})}{\text{argmax}} \langle \boldsymbol{\theta}, \boldsymbol{\mu} \rangle = \underset{\mathbf{y} \in \mathcal{Y}}{\text{argmax}} \langle \boldsymbol{\theta}, \mathbf{y} \rangle.$$

Why regularize outputs? The regularized prediction function (5) casts computing a prediction as a variational problem. It involves an optimization problem that balances between two terms: an “affinity” term $\langle \boldsymbol{\theta}, \boldsymbol{\mu} \rangle$, and a “confidence” term $\Omega(\boldsymbol{\mu})$ which should be low if $\boldsymbol{\mu}$ is “uncertain.” Two important classes of convex Ω are (squared) norms and, when $\text{dom}(\Omega)$ is the probability simplex, generalized negative entropies. However, our framework does **not** require Ω to be convex in general.

Introducing $\Omega(\boldsymbol{\mu})$ in (5) tends to move the prediction away from the vertices of $\text{conv}(\mathcal{Y})$. Unless the regularization term $\Omega(\boldsymbol{\mu})$ is negligible compared to the affinity term $\langle \boldsymbol{\theta}, \boldsymbol{\mu} \rangle$, a prediction becomes a convex combination of several vertices. As we shall see in §7, we can interpret this prediction as the mean under some underlying distribution. This contrasts with (4), which always outputs the most likely vertex, i.e., the mode.

Prediction over arbitrary domains. Regularized prediction functions are in fact not limited to convex hulls. We now state their precise definition in complete generality.

Definition 1 *Prediction function regularized by Ω*

Let $\Omega : \mathbb{R}^d \rightarrow \mathbb{R} \cup \{\infty\}$ be a regularization function, with $\text{dom}(\Omega) \subseteq \mathbb{R}^d$. The prediction function regularized by Ω is defined by

$$\hat{\mathbf{y}}_{\Omega}(\boldsymbol{\theta}) \in \underset{\boldsymbol{\mu} \in \text{dom}(\Omega)}{\text{argmax}} \langle \boldsymbol{\theta}, \boldsymbol{\mu} \rangle - \Omega(\boldsymbol{\mu}). \tag{5}$$

Allowing extended-real Ω permits general domain constraints in (5) via indicator functions. For instance, choosing $\Omega = I_{\text{conv}(\mathcal{Y})}$, where I_C is the indicator function defined in (1), recovers the MAP oracle (4). Importantly, the choice of domain $\text{dom}(\Omega)$ is not limited to convex hulls. For instance, we will also consider conic hulls, $\text{cone}(\mathcal{Y})$, later in this paper.

Choosing Ω . Regularized prediction functions $\hat{\mathbf{y}}_{\Omega}$ involve two main design choices: the domain $\text{dom}(\Omega)$ over which Ω is defined and Ω itself. The choice of $\text{dom}(\Omega)$ is mainly




Figure 1: **Illustration of the proposed regularized prediction framework over a convex hull $\text{conv}(\mathcal{Y})$.** A parametrized model $f_W: \mathcal{X} \rightarrow \mathbb{R}^d$ (linear model, neural network, etc.) produces a score vector $\theta \in \mathbb{R}^d$. The regularized prediction function \hat{y}_Ω produces a prediction $\hat{y} \in \text{conv}(\mathcal{Y})$. Regularized prediction functions are not limited to convex hulls and can be defined over arbitrary domains (Def. 1).

dictated by the type of output we want from \hat{y}_Ω , such as $\text{dom}(\Omega) = \text{conv}(\mathcal{Y})$ for convex combinations of elements of \mathcal{Y} , and $\text{dom}(\Omega) = \text{cone}(\mathcal{Y})$ for conic combinations. The choice of the regularization Ω itself further governs certain properties of \hat{y}_Ω , including, as we shall see in the sequel, its sparsity or its use of prior knowledge regarding the importance or misclassification cost of certain outputs. The choices of $\text{dom}(\Omega)$ and Ω may also be constrained by computational considerations. Indeed, while computing $\hat{y}_\Omega(\theta)$ involves a potentially challenging constrained maximization problem in general, we will see that certain choices of Ω lead to closed-form expressions. The power of our framework is that the user can focus solely on designing and computing \hat{y}_Ω : We will see in §3 how to automatically construct a loss function associated with \hat{y}_Ω .

2.2 Examples

To illustrate regularized prediction functions, we give several concrete examples enjoying a **closed-form expression**.

When $\Omega = I_{\Delta^d}$, $\hat{y}_\Omega(\theta)$ is a one-hot representation of the argmax prediction

$$\hat{y}_\Omega(\theta) \in \underset{p \in \Delta^d}{\text{argmax}} \langle \theta, p \rangle = \underset{y \in \{e_1, \dots, e_d\}}{\text{argmax}} \langle \theta, y \rangle.$$

We can see that output as a probability distribution that assigns all probability mass on the same class. When $\Omega = -\mathcal{H}^s + I_{\Delta^d}$, where $\mathcal{H}^s(p) := -\sum_i p_i \log p_i$ is Shannon’s entropy, $\hat{y}_\Omega(\theta)$ is the well-known softmax

$$\hat{y}_\Omega(\theta) = \text{softmax}(\theta) := \frac{\exp(\theta)}{\sum_{j=1}^d \exp(\theta_j)}. \quad (6)$$

See Boyd and Vandenberghe (2004, Ex. 3.25) for a derivation. The resulting distribution always has **dense** support. When $\Omega = \frac{1}{2} \|\cdot\|^2 + I_{\Delta^d}$, \hat{y}_Ω is the Euclidean projection onto

the probability simplex

$$\widehat{\mathbf{y}}_{\Omega}(\boldsymbol{\theta}) = \text{sparsemax}(\boldsymbol{\theta}) := \underset{\mathbf{p} \in \Delta^d}{\operatorname{argmin}} \|\mathbf{p} - \boldsymbol{\theta}\|^2,$$

a.k.a. the sparsemax transformation (Martins and Astudillo, 2016). It is well-known that

$$\underset{\mathbf{p} \in \Delta^d}{\operatorname{argmin}} \|\mathbf{p} - \boldsymbol{\theta}\|^2 = [\boldsymbol{\theta} - \tau \mathbf{1}]_+, \quad (7)$$

for some threshold $\tau \in \mathbb{R}$. Hence, the predicted distribution can have **sparse** support (it may assign exactly zero probability to low-scoring classes). The threshold τ can be computed **exactly** in $O(d)$ time (Brucker, 1984; Duchi et al., 2008; Condat, 2016).

The regularized prediction function paradigm is, however, not limited to the probability simplex: When $\Omega(\mathbf{p}) = -\sum_i \mathbb{H}^s([p_i, 1 - p_i]) + I_{[0,1]^d}(\mathbf{p})$, we get

$$\widehat{\mathbf{y}}_{\Omega}(\boldsymbol{\theta}) = \text{sigmoid}(\boldsymbol{\theta}) := \frac{\mathbf{1}}{\mathbf{1} + \exp(-\boldsymbol{\theta})},$$

i.e., the sigmoid function evaluated coordinate-wise. We can think of its output as a positive measure (unnormalized probability distribution).

We will see in §4 that the first three examples (argmax, softmax and sparsemax) are particular instances of a broader family of prediction functions, using the notion of **generalized entropy**. The last example is a special case of regularized prediction function over positive measures, developed in §6. Regularized prediction functions also encompass more complex convex polytopes for structured prediction, as we shall see in §7.

2.3 Gradient mapping and dual objective

From Danskin’s theorem (Danskin, 1966) (see also Bertsekas (1999, Proposition B.25)) $\widehat{\mathbf{y}}_{\Omega}(\boldsymbol{\theta})$ is a subgradient of Ω^* at $\boldsymbol{\theta}$, i.e., $\widehat{\mathbf{y}}_{\Omega}(\boldsymbol{\theta}) \in \partial\Omega^*(\boldsymbol{\theta})$. If, furthermore, Ω is strictly convex, then $\widehat{\mathbf{y}}_{\Omega}(\boldsymbol{\theta})$ is the gradient of Ω^* at $\boldsymbol{\theta}$, i.e., $\widehat{\mathbf{y}}_{\Omega}(\boldsymbol{\theta}) = \nabla\Omega^*(\boldsymbol{\theta})$. This interpretation of $\widehat{\mathbf{y}}_{\Omega}(\boldsymbol{\theta})$ as a (sub)gradient mapping will play a crucial role in the next section for deriving a loss function associated with $\widehat{\mathbf{y}}_{\Omega}(\boldsymbol{\theta})$.

Viewing $\widehat{\mathbf{y}}_{\Omega}(\boldsymbol{\theta})$ as the (sub)gradient of $\Omega^*(\boldsymbol{\theta})$ is also useful to derive the dual of (5). Let $\Omega := \Psi + \Phi$. It is well-known (Borwein and Lewis, 2010; Beck and Teboulle, 2012) that

$$\Omega^*(\boldsymbol{\theta}) = (\Psi + \Phi)^*(\boldsymbol{\theta}) = \inf_{\mathbf{u} \in \mathbb{R}^d} \Phi^*(\mathbf{u}) + \Psi^*(\boldsymbol{\theta} - \mathbf{u}) =: (\Phi^* \square \Psi^*)(\boldsymbol{\theta}), \quad (8)$$

where $f \square g$ denotes the infimal convolution of f with g . Furthermore, from Danskin’s theorem, $\widehat{\mathbf{y}}_{\Omega}(\boldsymbol{\theta}) = \nabla\Psi^*(\boldsymbol{\theta} - \mathbf{u}^*)$, where \mathbf{u}^* denotes an optimal solution of the infimum in (8). We can think of that infimum as the dual of the optimization problem in (5). When $\Psi = \Psi^* = \frac{1}{2}\|\cdot\|^2$, $\Omega^*(\boldsymbol{\theta})$ is known as the Moreau envelope of Φ^* (Moreau, 1965) and using Moreau’s decomposition, we obtain $\widehat{\mathbf{y}}_{\Omega}(\boldsymbol{\theta}) = \boldsymbol{\theta} - \text{prox}_{\Phi^*}(\boldsymbol{\theta}) = \text{prox}_{\Phi}(\boldsymbol{\theta})$. As another example, when $\Omega = \Psi + I_{\mathcal{C}}$, we obtain

$$\Omega^*(\boldsymbol{\theta}) = \inf_{\mathbf{u} \in \mathbb{R}^d} \sigma_{\mathcal{C}}(\mathbf{u}) + \Psi^*(\boldsymbol{\theta} - \mathbf{u}),$$




Figure 2: **Examples of regularized prediction functions.** The unregularized `argmax` prediction always hits a vertex of the probability simplex, leading to a probability distribution that puts all probability mass on the same class. Unlike, `softmax` which always occurs in the relative interior of the simplex and thus leads to a dense distribution, `sparsemax` (Euclidean projection onto the probability simplex) may hit the boundary, leading to a **sparse** probability distribution. We also display the `sigmoid` operator which lies in the unit cube and is thus not guaranteed to output a valid probability distribution.

where we used $I_{\mathcal{C}}^* = \sigma_{\mathcal{C}}$, the support function of \mathcal{C} . In particular, when $\mathcal{C} = \text{conv}(\mathcal{Y})$, we have $\sigma_{\mathcal{C}}(\mathbf{u}) = \max_{\mathbf{y} \in \mathcal{Y}} \langle \mathbf{u}, \mathbf{y} \rangle$. This dual view is informative insofar as it suggests that regularized prediction functions $\hat{\mathbf{y}}_{\Omega}(\boldsymbol{\theta})$ with $\Omega = \Psi + I_{\mathcal{C}}$ minimize a trade-off between maximizing the value achieved by the unregularized prediction function $\sigma_{\mathcal{C}}(\mathbf{u})$, and a proximity term $\Psi^*(\boldsymbol{\theta} - \mathbf{u})$.

2.4 Properties

We now discuss simple yet useful properties of regularized prediction functions. The first two assume that Ω is a symmetric function, i.e., that it satisfies

$$\Omega(\boldsymbol{\mu}) = \Omega(\mathbf{P}\boldsymbol{\mu}) \quad \forall \boldsymbol{\mu} \in \text{dom}(\Omega), \forall \mathbf{P} \in \mathcal{P},$$

where \mathcal{P} is the set of $d \times d$ permutation matrices.

Proposition 1 *Properties of regularized prediction functions $\hat{\mathbf{y}}_{\Omega}(\boldsymbol{\theta})$*

1. **Effect of a permutation.** *If Ω is symmetric, then $\forall \mathbf{P} \in \mathcal{P}: \hat{\mathbf{y}}_{\Omega}(\mathbf{P}\boldsymbol{\theta}) = \mathbf{P}\hat{\mathbf{y}}_{\Omega}(\boldsymbol{\theta})$.*
2. **Order preservation.** *Let $\boldsymbol{\mu} = \hat{\mathbf{y}}_{\Omega}(\boldsymbol{\theta})$. If Ω is symmetric, then the coordinates of $\boldsymbol{\mu}$ and $\boldsymbol{\theta}$ are sorted the same way, i.e., $\theta_i > \theta_j \Rightarrow \mu_i \geq \mu_j$ and $\mu_i > \mu_j \Rightarrow \theta_i > \theta_j$.*
3. **Approximation error.** *If Ω is γ -strongly convex and bounded with $L \leq \Omega(\boldsymbol{\mu}) \leq U$ for all $\boldsymbol{\mu} \in \text{dom}(\Omega)$, then $\frac{1}{2} \|\hat{\mathbf{y}}(\boldsymbol{\theta}) - \hat{\mathbf{y}}_{\Omega}(\boldsymbol{\theta})\|^2 \leq \frac{U-L}{\gamma}$.*

- 4. **Temperature scaling.** For any constant $t > 0$, $\widehat{\mathbf{y}}_{t\Omega}(\boldsymbol{\theta}) \in \partial\Omega^*(\boldsymbol{\theta}/t)$. If Ω is strictly convex, $\widehat{\mathbf{y}}_{t\Omega}(\boldsymbol{\theta}) = \widehat{\mathbf{y}}_{\Omega}(\boldsymbol{\theta}/t) = \nabla\Omega^*(\boldsymbol{\theta}/t)$.
- 5. **Constant invariance.** For any constant $c \in \mathbb{R}$, $\widehat{\mathbf{y}}_{\Omega+c}(\boldsymbol{\theta}) = \widehat{\mathbf{y}}_{\Omega}(\boldsymbol{\theta}; \mathbf{y})$.

The proof is given in Appendix B.1.

For classification, the order-preservation property ensures that the highest-scoring class according to $\boldsymbol{\theta}$ and $\widehat{\mathbf{y}}_{\Omega}(\boldsymbol{\theta})$ agree with each other:

$$\operatorname{argmax}_{i \in [d]} \theta_i = \operatorname{argmax}_{i \in [d]} (\widehat{\mathbf{y}}_{\Omega}(\boldsymbol{\theta}))_i.$$

Temperature scaling is useful to control how close we are to unregularized prediction functions. Clearly, $\widehat{\mathbf{y}}_{t\Omega}(\boldsymbol{\theta}) \rightarrow \widehat{\mathbf{y}}(\boldsymbol{\theta})$ as $t \rightarrow 0$, where $\widehat{\mathbf{y}}(\boldsymbol{\theta})$ is defined in (4).

3. Fenchel-Young losses

In the previous section, we introduced regularized prediction functions over arbitrary domains, as a generalization of classical (unregularized) decision functions. In this section, we introduce Fenchel-Young losses for learning models whose output layer is a regularized prediction function. We first give their definitions and state their properties (§3.1). We then discuss their relationship with Bregman divergences (§3.2) and their Bayes risk (§3.3). Finally, we show how to construct a cost-sensitive loss from any Fenchel-Young loss (§3.4).

3.1 Definition and properties

Given a regularized prediction function $\widehat{\mathbf{y}}_{\Omega}$, we define its associated loss as follows.

Definition 2 *Fenchel-Young loss generated by Ω*

Let $\Omega: \mathbb{R}^d \rightarrow \mathbb{R} \cup \{\infty\}$ be a regularization function such that the maximum in (5) is achieved for all $\boldsymbol{\theta} \in \mathbb{R}^d$. Let $\mathbf{y} \in \mathcal{Y} \subseteq \operatorname{dom}(\Omega)$ be a ground-truth label and $\boldsymbol{\theta} \in \operatorname{dom}(\Omega^*) = \mathbb{R}^d$ be a vector of prediction scores.

The **Fenchel-Young loss** $L_{\Omega}: \operatorname{dom}(\Omega^*) \times \operatorname{dom}(\Omega) \rightarrow \mathbb{R}_+$ generated by Ω is

$$L_{\Omega}(\boldsymbol{\theta}; \mathbf{y}) := \Omega^*(\boldsymbol{\theta}) + \Omega(\mathbf{y}) - \langle \boldsymbol{\theta}, \mathbf{y} \rangle. \tag{9}$$

It is easy to see that Fenchel-Young losses can be rewritten as

$$L_{\Omega}(\boldsymbol{\theta}; \mathbf{y}) = f_{\boldsymbol{\theta}}(\mathbf{y}) - f_{\boldsymbol{\theta}}(\widehat{\mathbf{y}}_{\Omega}(\boldsymbol{\theta})),$$

where $f_{\boldsymbol{\theta}}(\boldsymbol{\mu}) := \Omega(\boldsymbol{\mu}) - \langle \boldsymbol{\theta}, \boldsymbol{\mu} \rangle$, highlighting the relation with regularized prediction functions. Therefore, as long as we can compute a regularized prediction function $\widehat{\mathbf{y}}_{\Omega}(\boldsymbol{\theta})$, we can automatically obtain an associated Fenchel-Young loss $L_{\Omega}(\boldsymbol{\theta}; \mathbf{y})$. Conversely, we also have that $\widehat{\mathbf{y}}_{\Omega}$ outputs the prediction minimizing the loss:

$$\widehat{\mathbf{y}}_{\Omega}(\boldsymbol{\theta}) \in \operatorname{argmin}_{\boldsymbol{\mu} \in \operatorname{dom}(\Omega)} L_{\Omega}(\boldsymbol{\theta}; \boldsymbol{\mu}).$$

Table 1: **Examples of regularized prediction functions and their corresponding Fenchel-Young losses.** For multi-class classification, we assume $\mathcal{Y} = \{\mathbf{e}_i\}_{i=1}^d$ and the ground-truth is $\mathbf{y} = \mathbf{e}_k$, where \mathbf{e}_i denotes a standard basis (“one-hot”) vector. For structured classification, we assume that elements of \mathcal{Y} are d -dimensional binary vectors with $d \ll |\mathcal{Y}|$, and we denote by $\text{conv}(\mathcal{Y}) = \{\mathbb{E}_{\mathbf{p}}[Y] : \mathbf{p} \in \Delta^{|\mathcal{Y}|}\}$ the corresponding marginal polytope (Wainwright and Jordan, 2008). We denote by $\mathbf{H}^s(\mathbf{p}) := -\sum_i p_i \log p_i$ the Shannon entropy of a distribution $\mathbf{p} \in \Delta^{|\mathcal{Y}|}$.

Loss	$\text{dom}(\Omega)$	$\Omega(\boldsymbol{\mu})$	$\hat{\mathbf{y}}_{\Omega}(\boldsymbol{\theta})$	$L_{\Omega}(\boldsymbol{\theta}; \mathbf{y})$
Squared	\mathbb{R}^d	$\frac{1}{2}\ \boldsymbol{\mu}\ ^2$	$\boldsymbol{\theta}$	$\frac{1}{2}\ \mathbf{y} - \boldsymbol{\theta}\ ^2$
Perceptron	$\Delta^{ \mathcal{Y} }$	0	$\text{argmax}(\boldsymbol{\theta})$	$\max_i \theta_i - \theta_k$
Logistic	$\Delta^{ \mathcal{Y} }$	$-\mathbf{H}^s(\boldsymbol{\mu})$	$\text{softmax}(\boldsymbol{\theta})$	$\log \sum_i \exp \theta_i - \theta_k$
Hinge	$\Delta^{ \mathcal{Y} }$	$\langle \boldsymbol{\mu}, \mathbf{e}_k - \mathbf{1} \rangle$	$\text{argmax}(\mathbf{1} - \mathbf{e}_k + \boldsymbol{\theta})$	$\max_i [[i \neq k]] + \theta_i - \theta_k$
Sparsemax	$\Delta^{ \mathcal{Y} }$	$\frac{1}{2}\ \boldsymbol{\mu}\ ^2$	$\text{sparsemax}(\boldsymbol{\theta})$	$\frac{1}{2}\ \mathbf{y} - \boldsymbol{\theta}\ ^2 - \frac{1}{2}\ \hat{\mathbf{y}}_{\Omega}(\boldsymbol{\theta}) - \boldsymbol{\theta}\ ^2$
Logistic (one-vs-all)	$[0, 1]^{ \mathcal{Y} }$	$-\sum_i \mathbf{H}^s([\mu_i, 1 - \mu_i])$	$\text{sigmoid}(\boldsymbol{\theta})$	$\sum_i \log(1 + \exp(-(2y_i - 1)\theta_i))$
Structured perceptron	$\text{conv}(\mathcal{Y})$	0	$\text{MAP}(\boldsymbol{\theta})$	$\max_{\mathbf{y}'} \langle \boldsymbol{\theta}, \mathbf{y}' \rangle - \langle \boldsymbol{\theta}, \mathbf{y} \rangle$
Structured hinge	$\text{conv}(\mathcal{Y})$	$-\langle \boldsymbol{\mu}, \mathbf{c}_{\mathbf{y}} \rangle$	$\text{MAP}(\boldsymbol{\theta} + \mathbf{c}_{\mathbf{y}})$	$\max_{\mathbf{y}'} \langle \mathbf{c}_{\mathbf{y}}, \mathbf{y}' \rangle + \langle \boldsymbol{\theta}, \mathbf{y}' \rangle - \langle \boldsymbol{\theta}, \mathbf{y} \rangle$
CRF	$\text{conv}(\mathcal{Y})$	$\min_{\mathbf{p} \in \Delta^{ \mathcal{Y} } : \mathbb{E}_{\mathbf{p}}[Y] = \boldsymbol{\mu}} -\mathbf{H}^s(\mathbf{p})$	$\text{marginals}(\boldsymbol{\theta})$	$\log \sum_{\mathbf{y}'} \exp \langle \boldsymbol{\theta}, \mathbf{y}' \rangle - \langle \boldsymbol{\theta}, \mathbf{y} \rangle$
SparseMAP	$\text{conv}(\mathcal{Y})$	$\frac{1}{2}\ \boldsymbol{\mu}\ ^2$	$\text{SparseMAP}(\boldsymbol{\theta})$	$\frac{1}{2}\ \mathbf{y} - \boldsymbol{\theta}\ ^2 - \frac{1}{2}\ \hat{\mathbf{y}}_{\Omega}(\boldsymbol{\theta}) - \boldsymbol{\theta}\ ^2$

Examples of existing losses that fall into the Fenchel-Young loss family are given in Table 1. Some of these examples will be discussed in more details in the sequel of this paper.

Properties. As the name indicates, this family of loss functions is grounded in the Fenchel-Young inequality (Borwein and Lewis, 2010, Proposition 3.3.4)

$$\Omega^*(\boldsymbol{\theta}) + \Omega(\boldsymbol{\mu}) \geq \langle \boldsymbol{\theta}, \boldsymbol{\mu} \rangle \quad \forall \boldsymbol{\theta} \in \text{dom}(\Omega^*), \boldsymbol{\mu} \in \text{dom}(\Omega). \quad (10)$$

The inequality, together with well-known results regarding convex conjugates, imply the following properties of Fenchel-Young losses.

Proposition 2 *Properties of Fenchel-Young losses*

1. **Non-negativity.** $L_{\Omega}(\boldsymbol{\theta}; \mathbf{y}) \geq 0$ for any $\boldsymbol{\theta} \in \text{dom}(\Omega^*) = \mathbb{R}^d$ and $\mathbf{y} \in \mathcal{Y} \subseteq \text{dom}(\Omega)$.
2. **Zero loss.** If Ω is a lower semi-continuous proper convex function, then $\min_{\boldsymbol{\theta}} L_{\Omega}(\boldsymbol{\theta}; \mathbf{y}) = 0$, and $L_{\Omega}(\boldsymbol{\theta}; \mathbf{y}) = 0 \Leftrightarrow \mathbf{y} \in \partial\Omega^*(\boldsymbol{\theta})$. If Ω is strictly convex, then $L_{\Omega}(\boldsymbol{\theta}; \mathbf{y}) = 0 \Leftrightarrow \mathbf{y} = \hat{\mathbf{y}}_{\Omega}(\boldsymbol{\theta}) = \nabla\Omega^*(\boldsymbol{\theta}) = \text{argmin}_{\boldsymbol{\theta} \in \mathbb{R}^d} L_{\Omega}(\boldsymbol{\theta}; \mathbf{y})$.
3. **Convexity & subgradients.** L_{Ω} is convex in $\boldsymbol{\theta}$ and the residual vectors are its subgradients: $\hat{\mathbf{y}}_{\Omega}(\boldsymbol{\theta}) - \mathbf{y} \in \partial L_{\Omega}(\boldsymbol{\theta}; \mathbf{y})$.




Figure 3: **Illustration of the Fenchel-Young loss** $L_\Omega(\theta; \mathbf{y}) = \Omega^*(\theta) + \Omega(\mathbf{y}) - \langle \theta, \mathbf{y} \rangle$, here with $\Omega(\boldsymbol{\mu}) = \frac{1}{2}\|\boldsymbol{\mu}\|^2$ and $\text{dom}(\Omega) = \mathbb{R}^d$. Minimizing $L_\Omega(\theta; \mathbf{y})$ w.r.t. θ can be seen as minimizing the duality gap, the difference between $\Omega(\boldsymbol{\mu})$ and the tangent $\boldsymbol{\mu} \mapsto \langle \theta, \boldsymbol{\mu} \rangle - \Omega^*(\theta)$, at $\boldsymbol{\mu} = \mathbf{y}$ (the ground truth). The regularized prediction $\hat{\mathbf{y}}_\Omega(\theta)$ is the value of $\boldsymbol{\mu}$ at which the tangent touches $\Omega(\boldsymbol{\mu})$. When Ω is of Legendre type, $L_\Omega(\theta; \mathbf{y})$ is equal to the Bregman divergence generated by Ω between \mathbf{y} and $\hat{\mathbf{y}}_\Omega(\theta)$ (cf. §3.2). However, we do not require that assumption in this paper.

4. **Differentiability & smoothness.** *If Ω is strictly convex, then L_Ω is differentiable and $\nabla L_\Omega(\theta; \mathbf{y}) = \hat{\mathbf{y}}_\Omega(\theta) - \mathbf{y}$. If Ω is strongly convex, then L_Ω is smooth, i.e., $\nabla L_\Omega(\theta; \mathbf{y})$ is Lipschitz continuous.*
5. **Temperature scaling.** *For any constant $t > 0$, $L_{t\Omega}(\theta; \mathbf{y}) = tL_\Omega(\theta/t; \mathbf{y})$.*
6. **Constant invariance.** *For any constant $c \in \mathbb{R}$, $L_{\Omega+c}(\theta; \mathbf{y}) = L_\Omega(\theta; \mathbf{y})$.*

Remarkably, the non-negativity, convexity and constant invariance properties hold even if Ω is not convex. The zero loss property follows from the fact that, if Ω is l.s.c. proper convex, then (10) becomes an equality (i.e., the duality gap is zero) if and only if $\theta \in \partial\Omega(\boldsymbol{\mu})$. It suggests that the minimization of Fenchel-Young losses attempts to adjust the model to produce predictions $\hat{\mathbf{y}}_\Omega(\theta)$ that are close to the target \mathbf{y} , reducing the duality gap. This is illustrated with $\Omega = \frac{1}{2}\|\boldsymbol{\mu}\|^2$ and $\text{dom}(\Omega) = \mathbb{R}^d$ (leading to the squared loss) in Figure 3.

Domain of Ω^* . Our assumption that the maximum in the regularized prediction function (5) is achieved for all $\theta \in \mathbb{R}^d$ implies that $\text{dom}(\Omega^*) = \mathbb{R}^d$. This assumption is quite mild and does not require $\text{dom}(\Omega)$ to be bounded. Minimizing $L_\Omega(\theta; \mathbf{y})$ w.r.t. θ is therefore an **unconstrained** convex optimization problem. This contrasts with proper loss functions, which are defined over the probability simplex, as discussed in §10.

3.2 Relation with Bregman divergences

Fenchel-Young losses seamlessly work when $\mathcal{Y} = \text{dom}(\Omega)$ instead of $\mathcal{Y} \subset \text{dom}(\Omega)$. For example, in the case of the logistic loss, where $-\Omega$ is the Shannon entropy restricted to

Δ^d , allowing $\mathbf{y} \in \Delta^d$ instead of $\mathbf{y} \in \{\mathbf{e}_i\}_{i=1}^d$ yields the cross-entropy loss, $L_\Omega(\boldsymbol{\theta}; \mathbf{y}) = \text{KL}(\mathbf{y} \parallel \text{softmax}(\boldsymbol{\theta}))$, where KL denotes the (generalized) Kullback-Leibler divergence

$$\text{KL}(\mathbf{y} \parallel \boldsymbol{\mu}) := \sum_i y_i \log \frac{y_i}{\mu_i} - \sum_i y_i + \sum_i \mu_i.$$

This can be useful in a multi-label setting with supervision in the form of *label proportions*.

From this example, it is tempting to conjecture that a similar result holds for more general Bregman divergences (Bregman, 1967). Recall that the Bregman divergence $B_\Omega: \text{dom}(\Omega) \times \text{relint}(\text{dom}(\Omega)) \rightarrow \mathbb{R}_+$ generated by a strictly convex and differentiable Ω is

$$B_\Omega(\mathbf{y} \parallel \boldsymbol{\mu}) := \Omega(\mathbf{y}) - \Omega(\boldsymbol{\mu}) - \langle \nabla \Omega(\boldsymbol{\mu}), \mathbf{y} - \boldsymbol{\mu} \rangle. \quad (11)$$

In other words, this is the difference at \mathbf{y} between Ω and its linearization around $\boldsymbol{\mu}$. It turns out that $L_\Omega(\boldsymbol{\theta}; \mathbf{y})$ is not in general equal to $B_\Omega(\mathbf{y} \parallel \widehat{\mathbf{y}}_\Omega(\boldsymbol{\theta}))$. In fact the latter is **not necessarily convex** in $\boldsymbol{\theta}$ while the former always is. However, there is a duality relationship between Fenchel-Young losses and Bregman divergences, as we now discuss.

A “mixed-space” Bregman divergence. Letting $\boldsymbol{\theta} = \nabla \Omega(\mathbf{p})$ (i.e., $(\boldsymbol{\theta}, \mathbf{p})$ is a dual pair), we have $\Omega^*(\boldsymbol{\theta}) = \langle \boldsymbol{\theta}, \mathbf{p} \rangle - \Omega(\mathbf{p})$. Substituting in (11), we get $B_\Omega(\mathbf{y} \parallel \mathbf{p}) = L_\Omega(\boldsymbol{\theta}; \mathbf{y})$. In other words, Fenchel-Young losses can be viewed as a “mixed-form Bregman divergence” (Amari, 2016, Theorem 1.1) where the argument \mathbf{p} in (11) is **replaced by its dual point $\boldsymbol{\theta}$** . This difference is best seen by comparing the function signatures, $L_\Omega: \text{dom}(\Omega^*) \times \text{dom}(\Omega) \rightarrow \mathbb{R}_+$ vs. $B_\Omega: \text{dom}(\Omega) \times \text{relint}(\text{dom}(\Omega)) \rightarrow \mathbb{R}_+$. An important consequence is that Fenchel-Young losses do not impose any restriction on their left argument $\boldsymbol{\theta}$: Our assumption that the maximum in the prediction function (5) is achieved for all $\boldsymbol{\theta} \in \mathbb{R}^d$ implies $\text{dom}(\Omega^*) = \mathbb{R}^d$. In contrast, a Bregman divergence would typically need to be composed with a mapping from \mathbb{R}^d to $\text{dom}(\Omega)$, such as $\widehat{\mathbf{y}}_\Omega$, resulting in a possibly non-convex function.

Case of Legendre-type functions. We can make the relationship with Bregman divergences further precise when $\Omega = \Psi + I_C$, where Ψ is restricted to the class of so-called Legendre-type functions (Rockafellar, 1970; Wainwright and Jordan, 2008). We first recall the definition of this class of functions and then state our results.

Definition 3 *Essentially smooth and Legendre type functions*

A function Ψ is essentially smooth if

- $\text{dom}(\Psi)$ is non-empty,
- Ψ is differentiable throughout $\text{int}(\text{dom}(\Psi))$,
- and $\lim_{i \rightarrow \infty} \nabla \Psi(\boldsymbol{\mu}^i) = +\infty$ for any sequence $\{\boldsymbol{\mu}^i\}$ contained in $\text{dom}(\Psi)$, and converging to a boundary point of $\text{dom}(\Psi)$.

A function Ψ is of Legendre type if

- it is strictly convex on $\text{int}(\text{dom}(\Psi))$

- and essentially smooth.

For instance, $\Psi(\boldsymbol{\mu}) = \frac{1}{2}\|\boldsymbol{\mu}\|^2$ is Legendre-type with $\text{dom}(\Psi) = \mathbb{R}^d$, and $\Psi(\boldsymbol{\mu}) = \sum_i \mu_i \log \mu_i$ is Legendre-type with $\text{dom}(\Psi) = \mathbb{R}_+^d$. However, $\Omega(\boldsymbol{\mu}) = \frac{1}{2}\|\boldsymbol{\mu}\|^2 + I_{\mathbb{R}_+^d}(\boldsymbol{\mu})$ is not Legendre-type, since the gradient of Ω does not explode everywhere on the boundary of \mathbb{R}_+^d . The Legendre-type assumption crucially implies that

$$\nabla\Psi(\nabla\Psi^*(\boldsymbol{\theta})) = \boldsymbol{\theta} \quad \text{for all } \boldsymbol{\theta} \in \text{dom}(\Psi^*).$$

We can use this fact to derive the following results, proved in Appendix B.2.

Proposition 3 *Relation with Bregman divergences*

Let Ψ be of Legendre type with $\text{dom}(\Psi^*) = \mathbb{R}^d$ and let $\mathcal{C} \subseteq \text{dom}(\Psi)$ be a convex set. Let Ω be the restriction of Ψ to $\mathcal{C} \subseteq \text{dom}(\Psi)$, i.e., $\Omega := \Psi + I_{\mathcal{C}}$.

1. **Bregman projection.** The prediction function regularized by Ω , $\hat{\mathbf{y}}_{\Omega}(\boldsymbol{\theta})$, reduces to the Bregman projection of $\hat{\mathbf{y}}_{\Psi}(\boldsymbol{\theta})$ onto \mathcal{C} :

$$\hat{\mathbf{y}}_{\Omega}(\boldsymbol{\theta}) = \underset{\boldsymbol{\mu} \in \mathcal{C}}{\text{argmax}} \langle \boldsymbol{\theta}, \boldsymbol{\mu} \rangle - \Psi(\boldsymbol{\mu}) = \underset{\boldsymbol{\mu} \in \mathcal{C}}{\text{argmin}} B_{\Psi}(\boldsymbol{\mu} \| \hat{\mathbf{y}}_{\Psi}(\boldsymbol{\theta})).$$

2. **Difference of divergences.** For all $\boldsymbol{\theta} \in \mathbb{R}^d$ and $\mathbf{y} \in \mathcal{C}$:

$$L_{\Omega}(\boldsymbol{\theta}; \mathbf{y}) = B_{\Psi}(\mathbf{y} \| \hat{\mathbf{y}}_{\Psi}(\boldsymbol{\theta})) - B_{\Psi}(\hat{\mathbf{y}}_{\Omega}(\boldsymbol{\theta}) \| \hat{\mathbf{y}}_{\Psi}(\boldsymbol{\theta})).$$

3. **Bound.** For all $\boldsymbol{\theta} \in \mathbb{R}^d$ and $\mathbf{y} \in \mathcal{C}$:

$$0 \leq \underbrace{B_{\Psi}(\mathbf{y} \| \hat{\mathbf{y}}_{\Omega}(\boldsymbol{\theta}))}_{\text{possibly non-convex in } \boldsymbol{\theta}} \leq \underbrace{L_{\Omega}(\boldsymbol{\theta}; \mathbf{y})}_{\text{convex in } \boldsymbol{\theta}}$$

with equality when the loss is minimized

$$\hat{\mathbf{y}}_{\Omega}(\boldsymbol{\theta}) = \mathbf{y} \Leftrightarrow L_{\Omega}(\boldsymbol{\theta}; \mathbf{y}) = 0 \Leftrightarrow B_{\Psi}(\mathbf{y} \| \hat{\mathbf{y}}_{\Omega}(\boldsymbol{\theta})) = 0.$$

4. **Composite form.** When $\mathcal{C} = \text{dom}(\Psi)$, i.e., $\Omega = \Psi$, we have equality for all $\boldsymbol{\theta} \in \mathbb{R}^d$

$$L_{\Omega}(\boldsymbol{\theta}; \mathbf{y}) = B_{\Omega}(\mathbf{y} \| \hat{\mathbf{y}}_{\Omega}(\boldsymbol{\theta})).$$

We illustrate these properties using $\Psi = \frac{1}{2}\|\cdot\|^2$ as a running example. From the first property, since $\hat{\mathbf{y}}_{\Psi}(\boldsymbol{\theta}) = \boldsymbol{\theta}$ and $B_{\Psi}(\mathbf{y} \| \boldsymbol{\mu}) = \frac{1}{2}\|\mathbf{y} - \boldsymbol{\mu}\|^2$, we get

$$\hat{\mathbf{y}}_{\Omega}(\boldsymbol{\theta}) = \underset{\boldsymbol{\mu} \in \mathcal{C}}{\text{argmin}} B_{\Psi}(\boldsymbol{\mu} \| \hat{\mathbf{y}}_{\Psi}(\boldsymbol{\theta})) = \underset{\boldsymbol{\mu} \in \mathcal{C}}{\text{argmin}} B_{\Psi}(\boldsymbol{\mu} \| \boldsymbol{\theta}) = \underset{\boldsymbol{\mu} \in \mathcal{C}}{\text{argmin}} \|\boldsymbol{\mu} - \boldsymbol{\theta}\|^2,$$

recovering the Euclidean projection onto \mathcal{C} . The reduction of regularized prediction functions to Bregman projections (when Ψ is of Legendre type) is useful because there exist efficient algorithms for computing the Bregman projection onto various convex sets (Yasutake et al., 2011; Suehiro et al., 2012; Krichene et al., 2015; Lim and Wright, 2016). Therefore, we can use these algorithms to compute $\hat{\mathbf{y}}_{\Omega}(\boldsymbol{\theta})$ provided that $\hat{\mathbf{y}}_{\Psi}(\boldsymbol{\theta})$ is available.

From the second property, we obtain for all $\boldsymbol{\theta} \in \mathbb{R}^d$ and $\mathbf{y} \in \mathcal{C}$

$$L_\Omega(\boldsymbol{\theta}; \mathbf{y}) = \frac{1}{2} \|\mathbf{y} - \boldsymbol{\theta}\|^2 - \frac{1}{2} \|\widehat{\mathbf{y}}_\Omega(\boldsymbol{\theta}) - \boldsymbol{\theta}\|^2.$$

This recovers the expression of the sparsemax loss given in Table 1 with $\mathcal{C} = \Delta^d$.

From the third claim, we obtain for all $\boldsymbol{\theta} \in \mathbb{R}^d$ and $\mathbf{y} \in \mathcal{C}$

$$\frac{1}{2} \|\mathbf{y} - \widehat{\mathbf{y}}_\Omega(\boldsymbol{\theta})\|^2 \leq L_\Omega(\boldsymbol{\theta}; \mathbf{y}).$$

This shows that $L_\Omega(\boldsymbol{\theta}; \mathbf{y})$ provides a convex upper-bound for the possibly non-convex composite function $B_\Psi(\mathbf{y} \|\widehat{\mathbf{y}}_\Omega(\boldsymbol{\theta}))$. In particular, when $\mathcal{C} = \Delta^d$, we get $\frac{1}{2} \|\mathbf{y} - \text{sparsemax}(\boldsymbol{\theta})\|^2 \leq L_\Omega(\boldsymbol{\theta}; \mathbf{y})$. This suggests that the sparsemax loss is useful for sparse label proportion estimation, as confirmed in our experiments (§9).

Finally, from the last property, if $\Omega = \Psi = \frac{1}{2} \|\cdot\|^2$, we obtain $L_\Omega(\boldsymbol{\theta}; \mathbf{y}) = \frac{1}{2} \|\mathbf{y} - \boldsymbol{\theta}\|^2$, which is indeed the squared loss given in Table 1.

3.3 Expected loss, Bayes risk and Bregman information

In this section, we discuss the relation between the pointwise Bayes risk (minimal achievable loss) of a Fenchel-Young loss and Bregman information (Banerjee et al., 2005).

Expected loss. Let Y be a random variable taking values in \mathcal{Y} following the distribution $\mathbf{p} \in \Delta^{|\mathcal{Y}|}$. The expected loss (a.k.a. expected risk) is then

$$\begin{aligned} \mathbb{E}_{\mathbf{p}}[L_\Omega(\boldsymbol{\theta}; Y)] &= \sum_{\mathbf{y} \in \mathcal{Y}} p(\mathbf{y}) L_\Omega(\boldsymbol{\theta}; \mathbf{y}) \\ &= \sum_{\mathbf{y} \in \mathcal{Y}} p(\mathbf{y}) (\Omega^*(\boldsymbol{\theta}) + \Omega(\mathbf{y}) - \langle \boldsymbol{\theta}, \mathbf{y} \rangle) \\ &= \mathbb{E}_{\mathbf{p}}[\Omega(Y)] + \Omega^*(\boldsymbol{\theta}) - \langle \boldsymbol{\theta}, \mathbb{E}_{\mathbf{p}}[Y] \rangle \\ &= L_\Omega(\boldsymbol{\theta}; \mathbb{E}_{\mathbf{p}}[Y]) + \mathbb{I}_\Omega(Y; \mathbf{p}). \end{aligned} \tag{12}$$

Here, we defined the **Bregman information** of Y by

$$\mathbb{I}_\Omega(Y; \mathbf{p}) := \min_{\boldsymbol{\mu} \in \text{dom}(\Omega)} \mathbb{E}_{\mathbf{p}}[B_\Omega(Y \|\boldsymbol{\mu})] = \mathbb{E}_{\mathbf{p}}[B_\Omega(Y \|\mathbb{E}_{\mathbf{p}}[Y])] = \mathbb{E}_{\mathbf{p}}[\Omega(Y)] - \Omega(\mathbb{E}_{\mathbf{p}}[Y]).$$

We refer the reader to Banerjee et al. (2005) for a detailed discussion as to why the last two equalities hold. The r.h.s. is exactly equal to the difference between the two sides of Jensen's inequality $\mathbb{E}[\Omega(Y)] \geq \Omega(\mathbb{E}[Y])$ and is therefore non-negative. For this reason, it is sometimes also called Jensen gap (Reid and Williamson, 2011).

Bayes risk. From Proposition 2, we know that $\min_{\boldsymbol{\theta}} L_\Omega(\boldsymbol{\theta}; \mathbf{y}) = 0$ for all $\mathbf{y} \in \text{dom}(\Omega)$. Therefore, the pointwise Bayes risk coincides precisely with the Bregman information of Y ,

$$\min_{\boldsymbol{\theta} \in \mathbb{R}^d} \mathbb{E}_{\mathbf{p}}[L_\Omega(\boldsymbol{\theta}; Y)] = \min_{\boldsymbol{\theta} \in \mathbb{R}^d} L_\Omega(\boldsymbol{\theta}; \mathbb{E}_{\mathbf{p}}[Y]) + \mathbb{I}_\Omega(Y; \mathbf{p}) = \mathbb{I}_\Omega(Y; \mathbf{p}), \tag{13}$$

provided that $\mathbb{E}_{\mathbf{p}}[Y] \in \text{dom}(\Omega)$. A similar relation between Bayes risk and Bregman information exists for proper losses (Reid and Williamson, 2011). We can think of (13) as a measure of the “difficulty” of the task. Combining (12) and (13), we obtain

$$\mathbb{E}_{\mathbf{p}}[L_{\Omega}(\boldsymbol{\theta}; Y)] - \min_{\boldsymbol{\theta} \in \mathbb{R}^d} \mathbb{E}_{\mathbf{p}}[L_{\Omega}(\boldsymbol{\theta}; Y)] = L_{\Omega}(\boldsymbol{\theta}; \mathbb{E}_{\mathbf{p}}[Y]),$$

the pointwise “regret” of $\boldsymbol{\theta} \in \mathbb{R}^d$ w.r.t. $\mathbf{p} \in \Delta^{|\mathcal{Y}|}$. If $\mathcal{Y} = \{\mathbf{e}_i\}_{i=1}^d$, $L_{\Omega}(\boldsymbol{\theta}; \mathbb{E}_{\mathbf{p}}[Y]) = L_{\Omega}(\boldsymbol{\theta}; \mathbf{p})$.

3.4 Cost-sensitive losses

Fenchel-Young losses also include the hinge loss of support vector machines. Indeed, from any classification loss L_{Ω} , we can construct a cost-sensitive version of it as follows. Define

$$\Psi(\boldsymbol{\mu}; \mathbf{y}) := \Omega(\boldsymbol{\mu}) - \langle \mathbf{c}_{\mathbf{y}}, \boldsymbol{\mu} \rangle,$$

where $\mathbf{c}_{\mathbf{y}} \in \mathbb{R}_+^{|\mathcal{Y}|}$ is a fixed cost vector that may depend on the ground truth \mathbf{y} . For example, $\mathbf{c}_{\mathbf{y}} = \mathbf{1} - \mathbf{y}$ corresponds to the 0/1 cost and can be used to impose a margin. Then, L_{Ψ} is a cost-sensitive version of L_{Ω} , which can be written as

$$L_{\Psi(\cdot; \mathbf{y})}(\boldsymbol{\theta}; \mathbf{y}) = L_{\Omega}(\boldsymbol{\theta} + \mathbf{c}_{\mathbf{y}}; \mathbf{y}) = \Omega^*(\boldsymbol{\theta} + \mathbf{c}_{\mathbf{y}}) + \Omega(\mathbf{y}) - \langle \boldsymbol{\theta} + \mathbf{c}_{\mathbf{y}}, \mathbf{y} \rangle.$$

This construction recovers the multi-class hinge loss (Crammer and Singer (2001); $\Omega = 0$), the softmax-margin loss (Gimpel and Smith (2010); $\Omega = -\text{H}^{\text{S}}$), and the cost-augmented sparsemax (Shalev-Shwartz and Zhang (2016, Eq. (13)), Niculae et al. (2018); $\Omega = \frac{1}{2} \|\cdot\|^2$). It is easy to see that the associated regularized prediction function is

$$\hat{\mathbf{y}}_{\Psi}(\boldsymbol{\theta}) = \hat{\mathbf{y}}_{\Omega}(\boldsymbol{\theta} + \mathbf{c}_{\mathbf{y}}).$$

For the 0/1 cost $\mathbf{c}_{\mathbf{y}} = \mathbf{1} - \mathbf{y}$ and $\mathbf{y} = \mathbf{e}_k$, we have

$$\operatorname{argmax}_{i \in [d]} (\hat{\mathbf{y}}_{\Psi}(\boldsymbol{\theta}))_i = k \implies \operatorname{argmax}_{i \in [d]} (\hat{\mathbf{y}}_{\Omega}(\boldsymbol{\theta}))_i = k,$$

justifying the use of $\hat{\mathbf{y}}_{\Omega}$ at prediction time.

4. Probabilistic prediction with Fenchel-Young losses

In the previous section, we presented Fenchel-Young losses in a broad setting. We now restrict to classification over the probability simplex. More precisely, we restrict to the case $\mathcal{Y} = \{\mathbf{e}_i\}_{i=1}^d$, (i.e., unstructured multi-class classification), and assume that $\text{dom}(\Omega) \subseteq \text{conv}(\mathcal{Y}) = \Delta^d$. In this case, the regularized prediction function (5) becomes

$$\hat{\mathbf{y}}_{\Omega}(\boldsymbol{\theta}) \in \operatorname{argmax}_{\mathbf{p} \in \Delta^d} \langle \boldsymbol{\theta}, \mathbf{p} \rangle - \Omega(\mathbf{p}),$$

where $\boldsymbol{\theta} \in \mathbb{R}^d$ is a vector of (possibly negative) prediction scores produced by a model $\mathbf{f}_W(\mathbf{x})$ and $\mathbf{p} \in \Delta^d$ is a discrete probability distribution. It is a generalized exponential family

distribution (Grünwald and Dawid, 2004; Frongillo and Reid, 2014) with natural parameter $\theta \in \mathbb{R}^d$ and regularization Ω . Of particular interest is the case where $\hat{y}_\Omega(\theta)$ is **sparse**, meaning that there are scores θ for which the resulting $\hat{y}_\Omega(\theta)$ assigns zero probability to some classes. As seen in §2.2, this happens for example with the sparsemax transformation, but not with the softmax. Later, in §5, we will establish conditions for the regularized prediction function to be sparse and will connect it to the notion of separation margin.

We first discuss generalized entropies and the properties of the Fenchel-Young losses they induce (§4.1). We then discuss their expected loss, Bayes risk and Fisher consistency (§4.2). We then give examples of generalized entropies and corresponding loss functions, several of them new to our knowledge (§4.3). Finally, we discuss the binary classification setting, recovering several examples of commonly-used loss functions (§4.4).

4.1 Fenchel-Young loss generated by a generalized entropy

Generalized entropies. A natural choice of regularization function Ω over the probability simplex is $\Omega = -\mathsf{H}$, where H is a generalized entropy (Grünwald and Dawid, 2004), also called uncertainty function by DeGroot (1962): a concave function over Δ^d , used to measure the “uncertainty” in a distribution $\mathbf{p} \in \Delta^d$.

Assumptions: We will make the following assumptions about H .

- A.1.** Zero entropy: $\mathsf{H}(\mathbf{p}) = 0$ if \mathbf{p} is a delta distribution, i.e., $\mathbf{p} \in \{\mathbf{e}_i\}_{i=1}^d$.
- A.2.** Strict concavity: $\mathsf{H}((1 - \alpha)\mathbf{p} + \alpha\mathbf{p}') > (1 - \alpha)\mathsf{H}(\mathbf{p}) + \alpha\mathsf{H}(\mathbf{p}')$, for $\mathbf{p} \neq \mathbf{p}'$, $\alpha \in (0, 1)$.
- A.3.** Symmetry: $\mathsf{H}(\mathbf{p}) = \mathsf{H}(P\mathbf{p})$ for any $P \in \mathcal{P}$.

Assumptions A.2 and A.3 imply that H is Schur-concave (Bauschke and Combettes, 2017), a common requirement in generalized entropies. This in turn implies assumption A.1, up to a constant (that constant can easily be subtracted so as to satisfy assumption A.1). As suggested by the next result, proved in §B.3, together, these assumptions imply that H can be used as a sensible uncertainty measure.

Proposition 4 *If H satisfies assumptions A.1–A.3, then it is non-negative and uniquely maximized by the uniform distribution $\mathbf{p} = \mathbf{1}/d$.*

That is, assumptions A.1–A.3 ensure that the uniform distribution is the maximum entropy distribution. A particular case of generalized entropies satisfying assumptions A.1–A.3 are uniformly separable functions of the form $\mathsf{H}(\mathbf{p}) = \sum_{j=1}^d h(p_j)$, where $h : [0, 1] \rightarrow \mathbb{R}_+$ is a non-negative strictly concave function such that $h(0) = h(1) = 0$. However, our framework is not restricted to this form.

Induced Fenchel-Young loss. If the ground truth is $\mathbf{y} = \mathbf{e}_k$ and assumption A.1 holds, the Fenchel-Young loss definition (9) becomes

$$L_{-\mathsf{H}}(\theta; \mathbf{e}_k) = (-\mathsf{H})^*(\theta) - \theta_k. \quad (14)$$

This form was also recently proposed by [Duchi et al. \(2018, Proposition 3\)](#). By using the fact that $\Omega^*(\boldsymbol{\theta} + c\mathbf{1}) = \Omega^*(\boldsymbol{\theta}) + c$ for all $c \in \mathbb{R}$ if $\text{dom}(\Omega) \subseteq \Delta^d$, we can further rewrite it as

$$L_{-\text{H}}(\boldsymbol{\theta}; \mathbf{e}_k) = (-\text{H})^*(\boldsymbol{\theta} - \theta_k \mathbf{1}). \quad (15)$$

This expression shows that Fenchel-Young losses over Δ^d can be written solely in terms of the generalized ‘‘cumulant function’’ $(-\text{H})^*$. Indeed, when H is Shannon’s entropy, we recover the cumulant (a.k.a. log-partition) function $(-\text{H}^{\text{S}})^*(\boldsymbol{\theta}) = \log \sum_{i=1}^d \exp(\theta_i)$. When H is strongly concave over Δ^d , we can also see $(-\text{H})^*$ as a smoothed max operator ([Nesterov, 2005](#); [Niculae and Blondel, 2017](#); [Mensch and Blondel, 2018](#)) and hence $L_{-\text{H}}(\boldsymbol{\theta}; \mathbf{e}_k)$ can be seen as a smoothed upper-bound of the perceptron loss $(\boldsymbol{\theta}; \mathbf{e}_k) \mapsto \max_{i \in [d]} \theta_i - \theta_k$.

It is well-known that minimizing the logistic loss $(\boldsymbol{\theta}; \mathbf{e}_k) \mapsto \log \sum_i \exp \theta_i - \theta_k$ is equivalent to minimizing the KL divergence between \mathbf{e}_k and $\text{softmax}(\boldsymbol{\theta})$, which in turn is equivalent to maximizing the likelihood of the ground-truth label, $\text{softmax}(\boldsymbol{\theta})_k$. Importantly, this equivalence does not carry over for generalized entropies H : minimizing $L_{-\text{H}}(\boldsymbol{\theta}; \mathbf{e}_k)$ is not in general equivalent to minimizing $B_{-\text{H}}(\mathbf{e}_k, \hat{\boldsymbol{\gamma}}_{-\text{H}}(\boldsymbol{\theta}))$ or maximizing $\hat{\boldsymbol{\gamma}}_{-\text{H}}(\boldsymbol{\theta})_k$. In fact, maximizing the likelihood is generally a non-concave problem. Fenchel-Young losses can be seen as a principled way to construct a convex loss regardless of H .

4.2 Expected loss, Bayes risk and Fisher consistency

Let Y be a random variable taking values in $\mathcal{Y} = \{\mathbf{e}_i\}_{i=1}^d$. If H satisfies assumption A.1, we can use (14) to obtain simpler expressions of the expected loss and pointwise Bayes risk than the ones derived in §3.3.

Expected loss. Indeed, the expected loss (risk) for all $\mathbf{p} \in \Delta^d$ simplifies to:

$$\mathbb{E}_{\mathbf{p}}[L_{-\text{H}}(\boldsymbol{\theta}; Y)] = \sum_{i=1}^d p_i L_{-\text{H}}(\boldsymbol{\theta}; \mathbf{e}_i) = \sum_{i=1}^d p_i ((-\text{H})^*(\boldsymbol{\theta}) - \theta_i) = (-\text{H})^*(\boldsymbol{\theta}) - \langle \mathbf{p}, \boldsymbol{\theta} \rangle. \quad (16)$$

Bayes risk and entropy. The pointwise (conditional) Bayes risk thus becomes

$$\min_{\boldsymbol{\theta} \in \mathbb{R}^d} \mathbb{E}_{\mathbf{p}}[L_{-\text{H}}(\boldsymbol{\theta}; Y)] = \min_{\boldsymbol{\theta} \in \mathbb{R}^d} (-\text{H})^*(\boldsymbol{\theta}) - \langle \mathbf{p}, \boldsymbol{\theta} \rangle = \text{H}(\mathbf{p}). \quad (17)$$

Therefore, the pointwise Bayes risk is equal to the generalized entropy H generating $L_{-\text{H}}$, evaluated at \mathbf{p} . This is consistent with (13), which states that the pointwise Bayes risk is equal to the Bregman information of Y under \mathbf{p} , because

$$\mathbb{I}_{-\text{H}}(Y; \mathbf{p}) = -\mathbb{I}_{\text{H}}(Y; \mathbf{p}) = \text{H}(\mathbb{E}_{\mathbf{p}}[Y]) - \mathbb{E}_{\mathbf{p}}[\text{H}(Y)] = \text{H}(\mathbf{p}),$$

where we used $\mathbb{E}_{\mathbf{p}}[Y] = \mathbf{p}$ and assumption A.1.

A similar relation between generalized entropies and pointwise (conditional) Bayes risk is well-known in the proper loss (scoring rule) literature ([Grünwald and Dawid, 2004](#); [Gneiting and Raftery, 2007](#); [Reid and Williamson, 2010](#); [Williamson et al., 2016](#)). The main difference

is that the minimization above is over \mathbb{R}^d , while it is over Δ^d in that literature (§10.1). As noted by Reid and Williamson (2011, §4.6), Bregman information can also be connected to the notion of statistical information developed by DeGroot (1962), the reduction between prior and posterior uncertainty H , of which mutual information is a special case.

Fisher consistency. From (16), $\mathbb{E}_{\mathbf{p}}[L_{-\mathsf{H}}(\boldsymbol{\theta}; Y)] = L_{-\mathsf{H}}(\boldsymbol{\theta}; \mathbf{p}) + \Omega(\mathbf{p})$. Combined with Proposition 2, we have that the pointwise Bayes risk (17) is achieved if and only if $\hat{\boldsymbol{y}}_{-\mathsf{H}}(\boldsymbol{\theta}) = \mathbf{p}$. Such losses are Fisher consistent estimators of probabilities (Williamson et al., 2016).

4.3 Examples

We now give examples of generalized entropies over the simplex Δ^d (we omit the indicator function I_{Δ^d} from the definitions since there is no ambiguity). We illustrate them together with the regularized prediction function and loss they produce in Figure 4. Several of the resulting loss functions are new to our knowledge.

Shannon entropy (Shannon and Weaver, 1949). This is the foundation of information theory, defined as

$$\mathsf{H}^{\text{S}}(\mathbf{p}) := - \sum_{j=1}^d p_j \log p_j.$$

As seen in Table 1, the resulting Fenchel-Young loss $L_{-\mathsf{H}^{\text{S}}}$ corresponds to the logistic loss. The associated distribution is the classical softmax, Eq. (6).

Tsallis α -entropies (Tsallis, 1988). These entropies are defined as

$$\mathsf{H}_{\alpha}^{\text{T}}(\mathbf{p}) := k(\alpha - 1)^{-1} \left(1 - \sum_{j=1}^d p_j^{\alpha} \right),$$

where $\alpha \geq 0$ and k is an arbitrary positive constant. They arise as a generalization of the Shannon-Khinchin axioms to non-extensive systems (Suyari, 2004) and have numerous scientific applications (Gell-Mann and Tsallis, 2004; Martins et al., 2009a). For convenience, we set $k = \alpha^{-1}$ for the rest of this paper. Tsallis entropies satisfy assumptions A.1–A.3 and can also be written in separable form:

$$\mathsf{H}_{\alpha}^{\text{T}}(\mathbf{p}) := \sum_{j=1}^d h_{\alpha}(p_j) \quad \text{with} \quad h_{\alpha}(t) := \frac{t - t^{\alpha}}{\alpha(\alpha - 1)}. \quad (18)$$

The limit case $\alpha \rightarrow 1$ corresponds to the Shannon entropy. When $\alpha = 2$, we recover the Gini index (Gini, 1912), a popular “impurity measure” for decision trees:

$$\mathsf{H}_2^{\text{T}}(\mathbf{p}) = \frac{1}{2} \sum_{j=1}^d p_j(1 - p_j) = \frac{1}{2}(1 - \|\mathbf{p}\|_2^2) \quad \forall \mathbf{p} \in \Delta^d. \quad (19)$$

Using the constant invariance property in Proposition 2, it can be checked that $L_{-H_2^T}$ recovers the sparsemax loss (Martins and Astudillo, 2016) (cf. Table 1).

Another interesting case is $\alpha \rightarrow +\infty$, which gives $H_\infty^T(\mathbf{p}) = 0$, hence $L_{-H_\infty^T}$ is the perceptron loss in Table 1. The resulting “argmax” distribution puts all probability mass on the top-scoring classes. In summary, the prediction functions for $\alpha = 1, 2, \infty$ are respectively softmax, sparsemax, and argmax. Tsallis entropies can therefore be seen as a **continuous parametric family** subsuming these important cases.

Norm entropies. An interesting class of non-separable entropies are entropies generated by a q -norm, defined as

$$H_q^N(\mathbf{p}) := 1 - \|\mathbf{p}\|_q. \quad (20)$$

We call them **norm entropies**. By the Minkowski inequality, q -norms with $q > 1$ are strictly convex on the simplex, so H_q^N satisfies assumptions A.1–A.3 for $q > 1$. The resulting norm entropies differ from Tsallis entropies in that the norm is not raised to the power of q : a subtle but important difference. The limit case $q \rightarrow \infty$ is particularly interesting: in this case, we obtain $H_\infty^N = 1 - \|\cdot\|_\infty$, recovering the Berger-Parker dominance index (Berger and Parker, 1970), widely used in ecology to measure species diversity. We surprisingly encounter H_∞^N again in Section 5, as a limit case for the existence of separation margins.

Squared norm entropies. Inspired by Niculae and Blondel (2017), as a simple extension of the Gini index (19), we consider the generalized entropy based on squared q -norms:

$$H_q^{\text{sq}}(\mathbf{p}) := \frac{1}{2}(1 - \|\mathbf{p}\|_q^2) = \frac{1}{2} - \frac{1}{2} \left(\sum_{j=1}^d p_j^q \right)^{\frac{2}{q}}.$$

The constant term $\frac{1}{2}$, omitted by Niculae and Blondel (2017), ensures satisfaction of A.1. For $q \in (1, 2]$, it is known that the squared q -norm is strongly convex w.r.t. $\|\cdot\|_q$ (Ball et al., 1994), implying that $(-H_q^{\text{sq}})^*$, and therefore $L_{-H_q^{\text{sq}}}$, is smooth. Although $\hat{\mathbf{y}}_{-H_q^{\text{sq}}}(\boldsymbol{\theta})$ cannot to our knowledge be solved in closed form for $q \in (1, 2)$, efficient iterative algorithms such as projected gradient are available.

Rényi β -entropies. Rényi entropies (Rényi, 1961) are defined for any $\beta \geq 0$ as:

$$H_\beta^R(\mathbf{p}) := \frac{1}{1 - \beta} \log \sum_{j=1}^d p_j^\beta.$$

Unlike Shannon and Tsallis entropies, Rényi entropies are not separable, with the exception of $\beta \rightarrow 1$, which also recovers Shannon entropy as a limit case. The case $\beta \rightarrow +\infty$ gives $H_\beta^R(\mathbf{p}) = -\log \|\mathbf{p}\|_\infty$. For $\beta \in [0, 1]$, Rényi entropies satisfy assumptions A.1–A.3; for $\beta > 1$, Rényi entropies fail to be concave. They are however pseudo-concave (Mangasarian, 1965), meaning that, for all $\mathbf{p}, \mathbf{q} \in \Delta^d$, $\langle \nabla H_\beta^R(\mathbf{p}), \mathbf{q} - \mathbf{p} \rangle \leq 0$ implies $H_\beta^R(\mathbf{q}) \leq H_\beta^R(\mathbf{p})$. This implies, among other things, that points $\mathbf{p} \in \Delta^d$ with zero gradient are maximizers of $\langle \mathbf{p}, \boldsymbol{\theta} \rangle + H_\beta^R(\mathbf{p})$, which allows us to compute the predictive distribution $\hat{\mathbf{y}}_{-H_\beta^R}$ with gradient-based methods.




Figure 4: Examples of **generalized entropies** (left) along with their **prediction distribution** (middle) and **Fenchel-Young losses** (right) for the binary case, where $\mathbf{p} = [t, 1 - t] \in \Delta^2$ and $\boldsymbol{\theta} = [s, 0] \in \mathbb{R}^2$. Except for softmax, which never exactly reaches 0, all distributions shown on the center can have **sparse support**.

4.4 Binary classification case

When $d = 2$, the convex conjugate expression simplifies and reduces to a univariate maximization problem over $[0, 1]$:

$$\Omega^*(\boldsymbol{\theta}) = \max_{\mathbf{p} \in \Delta^2} \langle \boldsymbol{\theta}, \mathbf{p} \rangle - \Omega(\mathbf{p}) = \max_{p \in [0,1]} \theta_1 p + \theta_2(1-p) - \Omega([p, 1-p]) = \phi^*(\theta_1 - \theta_2) + \theta_2,$$

where we defined $\phi(p) := \Omega([p, 1-p])$. Likewise, it is easy to verify that we also have $\Omega^*(\boldsymbol{\theta}) = \phi^*(\theta_2 - \theta_1) + \theta_1$. Let us choose $\boldsymbol{\theta} = [s, -s] \in \mathbb{R}^2$ for some $s \in \mathbb{R}$. Since the ground truth is $\mathbf{y} \in \{\mathbf{e}_1, \mathbf{e}_2\}$, where \mathbf{e}_1 and \mathbf{e}_2 represent positive and negative classes, from (15), we can write

$$L_\Omega(\boldsymbol{\theta}; \mathbf{e}_j) = \Omega^*(\boldsymbol{\theta} - \theta_j \mathbf{1}) = \begin{cases} \Omega^*([0, -s]) = \phi^*(-s) & \text{if } j = 1 \\ \Omega^*([s, 0]) = \phi^*(s) & \text{if } j = 2 \end{cases}.$$

This can be written concisely as $\phi^*(-ys)$ where $y \in \{+1, -1\}$. Losses that can be written in this form are sometimes called margin losses (Reid and Williamson, 2010). These are losses that treat the positive and negative classes symmetrically. We have thus made a connection between margin losses and the regularization function $\phi(p) := \Omega([p, 1-p])$. Note, however, that this is different from the notion of margin we develop in §5.

Examples. Choosing $\phi(p) = -\mathbf{H}^s([p, 1-p])$ leads to the binary logistic loss

$$\phi^*(-ys) = \log(1 + \exp(-ys)).$$

Choosing $\phi(p) = -\mathbf{H}_2^T([p, 1-p]) = p^2 - p$ leads to

$$\phi^*(u) = \begin{cases} 0 & \text{if } u \leq -1 \\ u & \text{if } u \geq 1 \\ \frac{1}{4}(u+1)^2 & \text{o.w.} \end{cases}.$$

The resulting loss, $\phi(-ys)$, is known as the modified Huber loss in the literature (Zhang, 2004). Hence, the sparsemax loss is the multiclass extension of the modified Huber loss, as already noted in (Martins and Astudillo, 2016).

Finally, we note that the modified Huber loss is closely related to the smoothed hinge loss (Shalev-Shwartz and Zhang, 2016). Indeed, choosing $\phi(p) = \frac{1}{2}p^2 - p$ (notice the $\frac{1}{2}$ factor), we obtain

$$\phi^*(u) = \begin{cases} 0 & \text{if } u \leq -1 \\ u + \frac{1}{2} & \text{if } u \geq 0 \\ \frac{1}{2}(1+u)^2 & \text{o.w.} \end{cases}.$$

It can be verified that $\phi^*(-ys)$ is indeed the smoothed hinge loss.

5. Separation margin of Fenchel-Young losses

In this section, we are going to see that the simple assumptions A.1–A.3 about a generalized entropy H are enough to obtain results about the **separation margin** associated with $L_{-\mathsf{H}}$. The notion of margin is well-known in machine learning, lying at the heart of support vector machines and leading to generalization error bounds (Vapnik, 1998; Schölkopf and Smola, 2002; Guermeur, 2007). We provide a definition and will see that many other Fenchel-Young losses also have a “margin,” for suitable conditions on H . Then, we take a step further, and connect the existence of a margin with the **sparsity** of the regularized prediction function, providing necessary and sufficient conditions for Fenchel-Young losses to have a margin. Finally, we show how this margin can be computed analytically.

Definition 4 *Separation margin*

Let $L(\boldsymbol{\theta}; \mathbf{e}_k)$ be a loss function over $\mathbb{R}^d \times \{\mathbf{e}_i\}_{i=1}^d$. We say that L has the separation margin property if there exists $m > 0$ such that:

$$\theta_k \geq m + \max_{j \neq k} \theta_j \quad \Rightarrow \quad L(\boldsymbol{\theta}; \mathbf{e}_k) = 0. \quad (21)$$

The smallest possible m that satisfies (21) is called the margin of L , denoted $\text{margin}(L)$.

Examples. The most famous example of a loss with a separation margin is the **multi-class hinge loss**, $L(\boldsymbol{\theta}; \mathbf{e}_k) = \max\{0, \max_{j \neq k} 1 + \theta_j - \theta_k\}$, which we saw in Table 1 to be a Fenchel-Young loss: it is immediate from the definition that its margin is 1. Less trivially, Martins and Astudillo (2016, Prop. 3.5) showed that the **sparsemax loss** also has the separation margin property. On the negative side, the logistic loss does not have a margin, as it is strictly positive. Characterizing which Fenchel-Young losses have a margin is an important question which we address next.

Conditions for existence of margin. To accomplish our goal, we need to characterize the gradient mappings $\partial(-\mathsf{H})$ and $\nabla(-\mathsf{H})^*$ associated with generalized entropies (note that $\partial(-\mathsf{H})$ is never single-valued: if $\boldsymbol{\theta}$ is in $\partial(-\mathsf{H})(\mathbf{p})$, then so is $\boldsymbol{\theta} + c\mathbf{1}$, for any constant $c \in \mathbb{R}$). Of particular importance is the subdifferential set $\partial(-\mathsf{H})(\mathbf{e}_k)$. The next proposition, whose proof we defer to §B.4, uses this set to provide a necessary and sufficient condition for the existence of a separation margin, along with a formula for computing it.

Proposition 5 *Let H satisfy A.1–A.3. Then:*

1. The loss $L_{-\mathsf{H}}$ has a separation margin iff there is a $m > 0$ such that $m\mathbf{e}_k \in \partial(-\mathsf{H})(\mathbf{e}_k)$.
2. If the above holds, then the margin of $L_{-\mathsf{H}}$ is given by the smallest such m or, equivalently,

$$\text{margin}(L_{-\mathsf{H}}) = \sup_{\mathbf{p} \in \Delta^d} \frac{\mathsf{H}(\mathbf{p})}{1 - \|\mathbf{p}\|_\infty}. \quad (22)$$

Reassuringly, the first part confirms that the logistic loss does not have a margin, since $\partial(-\mathbf{H}^s)(\mathbf{e}_k) = \emptyset$. A second interesting fact is that the denominator of (22) is the generalized entropy $\mathbf{H}_\infty^N(\mathbf{p})$ introduced in §4: the **∞ -norm entropy**. As Figure 4 suggests, this entropy provides an upper bound for convex losses with unit margin. This provides some intuition to the formula (22), which seeks a distribution \mathbf{p} maximizing the **entropy ratio** between $\mathbf{H}(\mathbf{p})$ and $\mathbf{H}_\infty^N(\mathbf{p})$.

Relationship between sparsity and margins. The next result, proved in §B.5, characterizes more precisely the image of $\nabla(-\mathbf{H})^*$. In doing so, it establishes a key result in this paper: **a sufficient condition for the existence of a separation margin in $L_{-\mathbf{H}}$ is the sparsity of the regularized prediction function $\hat{\mathbf{y}}_{-\mathbf{H}} \equiv \nabla(-\mathbf{H})^*$, i.e., its ability to reach the entire simplex, including the boundary points.** If \mathbf{H} is uniformly separable, this is also a necessary condition.

Proposition 6 *Relationship between margin losses and sparse predictive probabilities*

Let \mathbf{H} satisfy A.1–A.3 and be uniformly separable, i.e., $\mathbf{H}(\mathbf{p}) = \sum_{i=1}^d h(p_i)$. Then the following statements are all equivalent:

1. $\partial(-\mathbf{H})(\mathbf{p}) \neq \emptyset$ for any $\mathbf{p} \in \Delta^d$;
2. The mapping $\nabla(-\mathbf{H})^*$ covers the full simplex, i.e., $\nabla(-\mathbf{H})^*(\mathbb{R}^d) = \Delta^d$;
3. $L_{-\mathbf{H}}$ has the separation margin property.

For a general \mathbf{H} (not necessarily separable) satisfying A.1–A.3, we have (1) \Leftrightarrow (2) \Rightarrow (3).

Let us reflect for a moment on the three conditions stated in Proposition 6. The first two conditions involve the subdifferential and gradient of $-\mathbf{H}$ and its conjugate; the third condition is the margin property of $L_{-\mathbf{H}}$. To provide some intuition, consider the case where \mathbf{H} is separable with $\mathbf{H}(\mathbf{p}) = \sum_i h(p_i)$ and h is differentiable in $(0, 1)$. Then, from the concavity of h , its derivative h' is decreasing, hence the first condition is met if $\lim_{t=0^+} h'(t) < \infty$ and $\lim_{t=1^-} h'(t) > -\infty$. This is the case with Tsallis entropies for $\alpha > 1$, but not Shannon entropy, since $h'(t) = -1 - \log t$ explodes at 0. As stated in Definition 3, functions whose gradient “explodes” in the boundary of their domain (hence failing to meet the first condition in Proposition 6) are called “essentially smooth” (Rockafellar, 1970). For those functions, $\nabla(-\mathbf{H})^*$ maps only to the relative interior of Δ^d , never attaining boundary points (Wainwright and Jordan, 2008); this is expressed in the second condition. This prevents essentially smooth functions from generating a sparse $\mathbf{y}_{-\mathbf{H}} \equiv \nabla(-\mathbf{H})^*$ or (if they are separable) a loss $L_{-\mathbf{H}}$ with a margin, as asserted by the third condition. Since Legendre-type functions (Definition 3) are strictly convex *and* essentially smooth, by Proposition 3, loss functions for which the composite form $L_{-\mathbf{H}}(\boldsymbol{\theta}; \mathbf{y}) = B_{-\mathbf{H}}(\mathbf{y} \|\hat{\mathbf{y}}_{-\mathbf{H}}(\boldsymbol{\theta}))$ holds, which is the case of the logistic loss but not of the sparsemax loss, do not enjoy a margin and cannot induce a sparse probability distribution. This is geometrically visible in Figure 4.

Margin computation. For Fenchel-Young losses that have the separation margin property, Proposition 5 provided a formula for determining the margin. While informative,

formula (22) is not very practical, as it involves a generally non-convex optimization problem. The next proposition, proved in §B.6, takes a step further and provides a remarkably simple closed-form expression for generalized entropies that are **twice-differentiable**. To simplify notation, we denote by $\nabla_j \mathbf{H}(\mathbf{p}) \equiv (\nabla \mathbf{H}(\mathbf{p}))_j$ the j^{th} component of $\nabla \mathbf{H}(\mathbf{p})$.

Proposition 7 *Assume \mathbf{H} satisfies the conditions in Proposition 6 and is twice-differentiable on the simplex. Then, for arbitrary $j \neq k$:*

$$\text{margin}(L_{-\mathbf{H}}) = \nabla_j \mathbf{H}(\mathbf{e}_k) - \nabla_k \mathbf{H}(\mathbf{e}_k). \quad (23)$$

In particular, if \mathbf{H} is separable, i.e., $\mathbf{H}(\mathbf{p}) = \sum_{i=1}^{|\mathcal{Y}|} h(p_i)$, where $h : [0, 1] \rightarrow \mathbb{R}_+$ is concave, twice differentiable, with $h(0) = h(1) = 0$, then

$$\text{margin}(L_{-\mathbf{H}}) = h'(0) - h'(1) = - \int_0^1 h''(t) dt. \quad (24)$$

The compact formula (23) provides a geometric characterization of separable entropies and their margins: (24) tells us that only the slopes of h at the two extremities of $[0, 1]$ are relevant in determining the margin.

Example: case of Tsallis and norm entropies. As seen in §4, Tsallis entropies are separable with $h(t) = (t - t^\alpha)/(\alpha(\alpha - 1))$. For $\alpha > 1$, $h'(t) = (1 - \alpha t^{\alpha-1})/(\alpha(\alpha - 1))$, hence $h'(0) = 1/(\alpha(\alpha - 1))$ and $h'(1) = -1/\alpha$. Proposition 7 then yields

$$\text{margin}(L_{-\mathbf{H}_\alpha^T}) = h'(0) - h'(1) = (\alpha - 1)^{-1}.$$

Norm entropies, while not separable, have gradient $\nabla \mathbf{H}_q^N(\mathbf{p}) = -(\mathbf{p}/\|\mathbf{p}\|_q)^{q-1}$, giving $\nabla \mathbf{H}_q^N(\mathbf{e}_k) = -\mathbf{e}_k$, so

$$\text{margin}(\mathbf{H}_q^N) = \nabla_j \mathbf{H}_q^N(\mathbf{e}_k) - \nabla_k \mathbf{H}_q^N(\mathbf{e}_k) = 1,$$

as confirmed visually in Figure 4, in the binary case.

6. Positive measure prediction with Fenchel-Young losses

In this section, we again restrict to classification and $\mathcal{Y} = \{\mathbf{e}_i\}_{i=1}^d$ but now assume that $\text{dom}(\Omega) \subseteq \text{cone}(\mathcal{Y}) = \mathbb{R}_+^d$, where $\text{cone}(\mathcal{Y})$ is the conic hull of \mathcal{Y} . In this case, the regularized prediction function (5) becomes

$$\hat{\mathbf{y}}_\Omega(\boldsymbol{\theta}) \in \underset{\mathbf{m} \in \mathbb{R}_+^d}{\text{argmax}} \langle \boldsymbol{\theta}, \mathbf{m} \rangle - \Omega(\mathbf{m}),$$

where $\boldsymbol{\theta} \in \mathbb{R}^d$ is again a vector of prediction scores and $\mathbf{m} \in \mathbb{R}_+^d$ can be interpreted as a discrete positive measure (unnormalized probability distribution).

We first demonstrate how Fenchel-Young losses over positive measures allow to recover **one-vs-all reductions** (§6.1), theoretically justifying this popular scheme. We then give examples of loss instantiations (§6.2).

6.1 Uniformly separable regularizers and one-vs-all loss functions

A particularly simple case is that of uniformly separable Ω , i.e., $\Omega(\mathbf{m}) = \sum_{j=1}^d \phi(m_j) + I_{\mathbb{R}_+^d}(\mathbf{m})$, for some $\phi: \mathbb{R}_+ \rightarrow \mathbb{R}$. In that case the regularized prediction function can be computed in a coordinate-wise fashion:

$$(\widehat{\mathbf{y}}_\Omega(\boldsymbol{\theta}))_j = \operatorname{argmax}_{m \in \mathbb{R}_+} m\theta_j - \phi(m).$$

As we shall later see, this simplified optimization problem often enjoys a closed-form solution. Intuitively, $(\widehat{\mathbf{y}}_\Omega(\boldsymbol{\theta}))_j$ can be interpreted as the “unnormalized probability” of class j . Likewise, the corresponding Fenchel-Young loss is separable over classes:

$$L_\Omega(\boldsymbol{\theta}; \mathbf{y}) = \Omega^*(\boldsymbol{\theta}) + \Omega(\mathbf{y}) - \langle \boldsymbol{\theta}, \mathbf{y} \rangle = \sum_{j=1}^d \phi^*(\theta_j) + \phi(y_j) - \theta_j y_j = \sum_{j=1}^d L_\phi(\theta_j; y_j),$$

where $y_j \in \{1, 0\}$ indicates membership to class j . The separability allows to train the model producing each θ_j in an embarrassingly parallel fashion.

Let us now consider the case $\phi(p) = -\mathbf{H}([p, 1 - p])$, where \mathbf{H} is a generalized entropy and ϕ is restricted to $[0, 1]$. From §4.4, we know that $\phi^*(s) = \phi^*(-s) + s$ for all $s \in \mathbb{R}$. If \mathbf{H} further satisfies assumption A.1, meaning that $\phi(1) = \phi(0) = 0$, then we obtain

$$L_\phi(s; 1) = \phi^*(s) + \phi(1) - s = \phi^*(-s) \quad \text{and} \quad L_\phi(s; 0) = \phi^*(s) + \phi(0) = \phi^*(s).$$

Combining the two, we obtain

$$L_\Omega(\boldsymbol{\theta}; \mathbf{y}) = \sum_{j=1}^d \phi^*(-(2y_j - 1)\theta_j),$$

where we used that $(2y_j - 1) \in \{+1, -1\}$. Fenchel-Young losses thus recover classical **one-vs-all** loss functions. Since Fenchel-Young losses satisfy $L_\Omega(\boldsymbol{\theta}; \mathbf{y}) = 0 \Leftrightarrow \widehat{\mathbf{y}}_\Omega(\boldsymbol{\theta}) = \mathbf{y}$, our framework provides a theoretical justification for one-vs-all, a scheme that works well in practice (Rifkin and Klautau, 2004). Further, our framework justifies using $\widehat{\mathbf{y}}_\Omega(\boldsymbol{\theta})$ as a measure of class membership, as commonly implemented (with post-normalization) in software packages (Pedregosa et al., 2011; Buitinck et al., 2013).

Finally, we point out that in the binary case, choosing $\boldsymbol{\theta} = [s, -s]$, we have the relationship

$$L_\Omega(\boldsymbol{\theta}; \mathbf{e}_1) = 2\phi^*(-s) \quad \text{and} \quad L_\Omega(\boldsymbol{\theta}; \mathbf{e}_2) = 2\phi^*(s).$$

Thus, we recover the same loss as in §4.4, up to a constant 2 factor (i.e., learning over the simplex Δ^d or over the unit cube $[0, 1]^d$ is equivalent for binary classification).

6.2 Examples

Recall that $\Omega(\mathbf{m}) = \sum_{j=1}^d \phi(m_j) + I_{\mathbb{R}_+^d}(\mathbf{m})$. Choosing $\phi(p) = \frac{1}{2}p^2$ gives

$$\widehat{\mathbf{y}}_\Omega(\boldsymbol{\theta}) = [\boldsymbol{\theta}]_+ \quad \text{and} \quad L_\Omega(\boldsymbol{\theta}; \mathbf{y}) = \sum_{i=1}^d \frac{1}{2}[\theta_i]_+^2 + \frac{1}{2}y_i^2 - \theta_i y_i.$$

Choosing $\phi(p) = -\mathbf{H}^s([p, 1-p]) + I_{[0,1]}(p)$ leads to

$$\widehat{\mathbf{y}}_\Omega(\boldsymbol{\theta}) = \text{sigmoid}(\boldsymbol{\theta}) := \frac{\mathbf{1}}{\mathbf{1} + \exp(-\boldsymbol{\theta})}, \quad \text{and} \quad L_\Omega(\boldsymbol{\theta}; \mathbf{y}) = \sum_{i=1}^d \log(1 + \exp(-(2y_i - 1)\theta_i))$$

the classical sigmoid function and the one-vs-all logistic function. Note that $-\Omega(\mathbf{p}) = \sum_j \mathbf{H}^s([p_j, 1-p_j]) = -\sum_j p_j \log p_j + (1-p_j) \log(1-p_j)$ is sometimes known as the Fermi-Dirac entropy (Borwein and Lewis, 2010, Commentary of Section 3.3).

Choosing $\phi(p) = -\mathbf{H}_2^s([p, 1-p]) + I_{[0,1]}(p) = p^2 - p + I_{[0,1]}(p)$ leads to

$$(\widehat{\mathbf{y}}_\Omega(\boldsymbol{\theta}))_j = (\phi^*)'(s_j) = \begin{cases} 0 & \text{if } u \leq -1 \\ 1 & \text{if } u \geq 1 \\ \frac{1}{2}(u+1) & \text{o.w.} \end{cases} \quad \forall j \in [d].$$

We can think of this function as a “sparse sigmoid”.

7. Structured prediction with Fenchel-Young losses

In this section, we now turn to prediction over the convex hull of a set \mathcal{Y} of structured objects (sequences, trees, assignments, etc.), represented as d -dimensional vectors, i.e. $\mathcal{Y} \subseteq \mathbb{R}^d$. In this case, the regularized prediction function (5) becomes

$$\widehat{\mathbf{y}}_\Omega(\boldsymbol{\theta}) = \underset{\boldsymbol{\mu} \in \text{conv}(\mathcal{Y})}{\text{argmax}} \langle \boldsymbol{\mu}, \boldsymbol{\theta} \rangle - \Omega(\boldsymbol{\mu}).$$

Typically, $|\mathcal{Y}|$ will be exponential in d , i.e., $d \ll |\mathcal{Y}|$. Because $\widehat{\mathbf{y}}_\Omega(\boldsymbol{\theta}) = \sum_{\mathbf{y} \in \mathcal{Y}} p(\mathbf{y}) \mathbf{y} = \mathbb{E}_{\mathbf{p}}[Y]$ for some $\mathbf{p} \in \Delta^{|\mathcal{Y}|}$, $\widehat{\mathbf{y}}_\Omega(\boldsymbol{\theta})$ can be interpreted as the expectation under some (not necessarily unique) underlying distribution.

We first discuss the concepts of probability distribution regularization (§7.1) and mean regularization (§7.2), and the implications in terms of computational tractability and identifiability (note that in the unstructured setting, probability distribution and mean regularizations coincide). We then give several examples of polytopes and show how a Fenchel-Young loss can seamlessly be constructed over them given access to regularized prediction functions (§7.3). Finally, we extend the notion of separation margin to the structured setting (§7.4).

7.1 Distribution regularization, marginal inference and structured sparsemax

In this section, we discuss the concept of **probability distribution regularization**. Our treatment follows closely the variational formulations of exponential families (Barndorff-Nielsen, 1978; Wainwright and Jordan, 2008) and generalized exponential families (Grünwald and Dawid, 2004; Frongillo and Reid, 2014) but adopts the novel viewpoint of regularized prediction functions. We discuss two instances of that framework: the structured counterpart of the softmax, marginal inference, and a new structured counterpart of the sparsemax, structured sparsemax.

Let us start with a probability-space regularized prediction function

$$\hat{\mathbf{y}}_{-\mathbf{H}}(\mathbf{s}_\theta) = \operatorname{argmax}_{\mathbf{p} \in \Delta^{|\mathcal{Y}|}} \langle \mathbf{p}, \mathbf{s}_\theta \rangle + \mathbf{H}(\mathbf{p}),$$

where \mathbf{H} is a generalized entropy and $\mathbf{s}_\theta := (\langle \theta, \mathbf{y} \rangle)_{\mathbf{y} \in \mathcal{Y}} \in \mathbb{R}^{|\mathcal{Y}|}$ is a vector that contains the scores of all possible structures. The above optimization problem is defined over a space of size $|\mathcal{Y}|$ and is thus intractable in the general case. The regularized prediction function outputs a probability distribution over structures and from Danskin's theorem, we have

$$\hat{\mathbf{y}}_{-\mathbf{H}}(\mathbf{s}_\theta) = \nabla_{\mathbf{s}_\theta} (-\mathbf{H})^*(\mathbf{s}_\theta) = \mathbf{p}^*.$$

Using the chain rule, differentiating w.r.t. θ instead of \mathbf{s}_θ gives

$$\nabla_{\theta} (-\mathbf{H})^*(\mathbf{s}_\theta) = \mathbb{E}_{\mathbf{p}^*}[Y] \in \operatorname{conv}(\mathcal{Y}).$$

Therefore, the gradient w.r.t. θ corresponds to the mean under \mathbf{p}^* . This can be equivalently expressed under our framework by defining a regularization function directly in mean space. Let us define the regularization function Ω over $\operatorname{conv}(\mathcal{Y})$ as

$$-\Omega(\boldsymbol{\mu}) := \max_{\mathbf{p} \in \Delta^{|\mathcal{Y}|}} \mathbf{H}(\mathbf{p}) \quad \text{s.t.} \quad \mathbb{E}_{\mathbf{p}}[Y] = \boldsymbol{\mu}. \quad (25)$$

That is, among all distributions satisfying the first-moment matching condition $\mathbb{E}_{\mathbf{p}}[Y] = \boldsymbol{\mu}$, we seek the distribution \mathbf{p}^* with maximum (generalized) entropy. This allows to make the underlying distribution unique and identifiable. With this choice of Ω , simple calculations show that the corresponding regularized prediction function is precisely the mean under that distribution:

$$\hat{\mathbf{y}}_{\Omega}(\theta) = \nabla^* \Omega(\theta) = \mathbb{E}_{\mathbf{p}^*}[Y] \in \operatorname{conv}(\mathcal{Y}).$$

Similar results hold for higher order moments: If Ω^* is twice-differentiable, the Hessian corresponds to the second moment under \mathbf{p}^* :

$$\operatorname{cov}_{\mathbf{p}^*}[Y] = \nabla^2 \Omega^*(\theta).$$

The relation between these various maps is summarized in Figure 5.

Marginal inference. A particular case of (25) is

$$-\Omega(\boldsymbol{\mu}) = \max_{\mathbf{p} \in \Delta^{|\mathcal{Y}|}} \mathbf{H}^s(\mathbf{p}) \quad \text{s.t.} \quad \mathbb{E}_{\mathbf{p}}[Y] = \boldsymbol{\mu}, \quad (26)$$

where \mathbf{H}^s is the Shannon entropy. The distribution achieving maximum Shannon entropy, which is unique, is the Gibbs distribution $p(\mathbf{y}; \theta) \propto \exp(\langle \theta, \mathbf{y} \rangle)$.

As shown in Table 1, the resulting loss L_{Ω} is the CRF loss (Lafferty et al., 2001) and the resulting regularized prediction function $\hat{\mathbf{y}}_{\Omega}$ is known as marginal inference in the literature (Wainwright and Jordan, 2008):

$$\hat{\mathbf{y}}_{\Omega}(\theta) = \operatorname{marginals}(\theta) := \sum_{\mathbf{y} \in \mathcal{Y}} \exp(\langle \theta, \mathbf{y} \rangle) \mathbf{y} / Z(\theta),$$

where $Z(\theta) := \sum_{\mathbf{y} \in \mathcal{Y}} \exp(\langle \theta, \mathbf{y} \rangle)$ is the partition function (normalization constant) of the Gibbs distribution. Although marginal inference is intractable in general, we will see in §7.3 that it can be computed exactly and efficiently for specific polytopes. The conjugate of $\Omega(\boldsymbol{\mu})$, $\Omega^*(\theta)$, corresponds to the log partition function: $\Omega^*(\theta) = \log Z(\theta)$.




Figure 5: **Summary of maps between spaces.** We define Ω over $\text{conv}(\mathcal{Y})$ using a generalized maximum entropy principle: $-\Omega(\boldsymbol{\mu}) := \max_{\mathbf{p} \in \Delta^{|\mathcal{Y}|}} \mathbf{H}(\mathbf{p})$ s.t. $\mathbb{E}_{\mathbf{p}}[Y] = \boldsymbol{\mu}$ (cf. §7.1). We also define $\mathbf{s}_{\boldsymbol{\theta}} := (\langle \boldsymbol{\theta}, \mathbf{y} \rangle)_{\mathbf{y} \in \mathcal{Y}} \in \mathbb{R}^{|\mathcal{Y}|}$, the vector that contains the scores of all possible structures, a linear map from \mathbb{R}^d to $\mathbb{R}^{|\mathcal{Y}|}$. Note that the diagram is not necessarily commutative.

Structured sparsemax. From the above perspective, it is tempting to define a sparsemax counterpart of the CRF loss and of marginal inference by replacing Shannon’s entropy \mathbf{H}^S with the Gini index \mathbf{H}_2^T (see (19) for a definition) in (26):

$$-\Omega(\boldsymbol{\mu}) = \max_{\mathbf{p} \in \Delta^{|\mathcal{Y}|}} \mathbf{H}_2^T(\mathbf{p}) \quad \text{s.t.} \quad \mathbb{E}_{\mathbf{p}}[Y] = \boldsymbol{\mu}.$$

This is equivalent to seeking a distribution \mathbf{p}^* with minimum squared norm

$$\Omega(\boldsymbol{\mu}) + \frac{1}{2} = \min_{\mathbf{p} \in \Delta^{|\mathcal{Y}|}} \frac{1}{2} \|\mathbf{p}\|^2 \quad \text{s.t.} \quad \mathbb{E}_{\mathbf{p}}[Y] = \boldsymbol{\mu}.$$

The corresponding $\hat{\mathbf{y}}_{\Omega}(\boldsymbol{\theta})$ is then $\hat{\mathbf{y}}_{\Omega}(\boldsymbol{\theta}) = \mathbb{E}_{\mathbf{p}^*}[Y]$. Recalling that $\mathbf{p}^* = \hat{\mathbf{y}}_{-\mathbf{H}_2^T}(\mathbf{s}_{\boldsymbol{\theta}})$, a naive approach would be to first compute the score vector $\mathbf{s}_{\boldsymbol{\theta}} = (\langle \boldsymbol{\theta}, \mathbf{y} \rangle)_{\mathbf{y} \in \mathcal{Y}} \in \mathbb{R}^{|\mathcal{Y}|}$ and then to project that vector on the simplex to obtain \mathbf{p}^* . Unfortunately, $\mathbf{s}_{\boldsymbol{\theta}}$ could be exponentially large and \mathbf{p}^* could be arbitrarily dense in the worst case, making that approach intractable in general. As also noted by Pillutla et al. (2018), we can approximate the problem by further imposing an upper-bound on the sparsity of the distribution, $\|\mathbf{p}\|_0 \leq k$. Since the optimal k -sparse distribution solely depends on the top- k elements of $\mathbf{s}_{\boldsymbol{\theta}}$ (Kyrillidis et al., 2013), we can first use a k -best oracle to retrieve the top- k elements of $\mathbf{s}_{\boldsymbol{\theta}}$ and project them onto the simplex. A disadvantage of that approach is that k -best oracles are usually more complex than MAP oracles and k could be arbitrarily large in order to guarantee an exact solution to the original problem. For shortest path problems over a directed acyclic graph, an alternative approximation is to smooth the maximum operator directly in the dynamic programming recursion (Mensch and Blondel, 2018).

7.2 Mean regularization and SparseMAP

Marginal inference is computationally tractable only for certain polytopes and structured sparsemax requires approximations, even for polytopes for which exact marginal inference

is available. In this section, we discuss an alternative approach which only requires access to a MAP oracle, broadening the set of applicable polytopes. The key idea is **mean regularization**: We directly regularize the mean $\boldsymbol{\mu}$ rather than the distribution \boldsymbol{p} .

The mean regularization counterpart of marginal inference is

$$\hat{\boldsymbol{y}}_{-H^S}(\boldsymbol{\theta}) = \operatorname{argmax}_{\boldsymbol{\mu} \in \operatorname{conv}(\mathcal{Y})} \langle \boldsymbol{\theta}, \boldsymbol{\mu} \rangle + H^S(\boldsymbol{\mu}) = \operatorname{argmin}_{\boldsymbol{\mu} \in \operatorname{conv}(\mathcal{Y})} \operatorname{KL}(\boldsymbol{\mu} \| e^{\boldsymbol{\theta}-1}).$$

It has been used for specific convex polytopes, most importantly in the optimal transport literature (Cuturi, 2013; Peyré and Cuturi, 2017) but also for learning to predict permutation matrices (Helmhold and Warmuth, 2009) or permutations (Yasutake et al., 2011; Ailon et al., 2016). The mean regularization counterpart of sparsemax is known as SparseMAP (Nicolae et al., 2018):

$$\hat{\boldsymbol{y}}_{\Omega}(\boldsymbol{\theta}) = \operatorname{SparseMAP}(\boldsymbol{\theta}) := \operatorname{argmax}_{\boldsymbol{\mu} \in \operatorname{conv}(\mathcal{Y})} \langle \boldsymbol{\theta}, \boldsymbol{\mu} \rangle - \frac{1}{2} \|\boldsymbol{\mu}\|^2 = \operatorname{argmin}_{\boldsymbol{\mu} \in \operatorname{conv}(\mathcal{Y})} \|\boldsymbol{\mu} - \boldsymbol{\theta}\|^2.$$

The main advantage of mean regularization is computational. Indeed, computing $\hat{\boldsymbol{y}}_{\Omega}(\boldsymbol{\theta})$ now simply involves a d -dimensional optimization problem instead of a $|\mathcal{Y}|$ -dimensional one and can be cast as a Bregman projection onto $\operatorname{conv}(\mathcal{Y})$ (§3.2). For specific polytopes, that projection can often be computed directly (§7.3). More generally, it can always be computed to arbitrary precision given access to a MAP oracle

$$\operatorname{argmax}_{\boldsymbol{y} \in \operatorname{conv}(\mathcal{Y})} \langle \boldsymbol{\theta}, \boldsymbol{y} \rangle = \operatorname{argmax}_{\boldsymbol{y} \in \mathcal{Y}} \langle \boldsymbol{\theta}, \boldsymbol{y} \rangle =: \operatorname{MAP}(\boldsymbol{\theta}),$$

thanks to conditional gradient algorithms (§8.3). Since MAP inference is a cornerstone of structured prediction, efficient algorithms have been developed for many kinds of structures (§7.3). In addition, as conditional gradient algorithms maintain a convex combination of vertices, they can also return a (not necessarily unique) distribution $\boldsymbol{p} \in \Delta^{|\mathcal{Y}|}$ such that $\hat{\boldsymbol{y}}_{\Omega}(\boldsymbol{\theta}) = \mathbb{E}_{\boldsymbol{p}}[Y]$. From Carathéodory’s theorem, the support of \boldsymbol{p} contains at most $d \ll |\mathcal{Y}|$ elements. This ensures that $\hat{\boldsymbol{y}}_{\Omega}(\boldsymbol{\theta})$ can be written as a “small” number of elementary structures. The price to pay for this computational tractability is that the underlying distribution \boldsymbol{p} is not necessarily unique.

7.3 Examples

Deriving loss functions for structured outputs can be challenging. In this section, we give several examples of polytopes $\operatorname{conv}(\mathcal{Y})$ for which efficient computational oracles (MAP, marginal inference, projection) are available, thus allowing to obtain a Fenchel-Young loss for learning over these polytopes. A summary is given in Table 7.3.

Sequences. We wish to tag a sequence $(\boldsymbol{x}_1, \dots, \boldsymbol{x}_n)$ of vectors in \mathbb{R}^p (e.g., word representations) with the most likely output sequence (e.g., entity tags) $\boldsymbol{s} = (s_1, \dots, s_n) \in [m]^n$. It is convenient to represent each sequence \boldsymbol{s} as a $n \times m \times m$ binary tensor $\boldsymbol{y} \in \mathcal{Y}$, such that $y_{t,i,j} = 1$ if \boldsymbol{y} transitions from node j to node i on time t , and 0 otherwise. The potentials $\boldsymbol{\theta}$

Table 2: **Examples of convex polytopes and computational cost of the three main computational oracles:** MAP, marginal inference and projection (in the Euclidean distance sense or more generally in the Bregman divergence sense). For polytopes for which direct marginal inference or projection algorithms are not available, we can always use conditional gradient algorithms to compute an optimal solution to arbitrary precision — see §8.3.

Polytope	Vertices	Dim.	MAP	Marginal	Projection
Probability simplex	Basis vectors	d	$O(d)$	$O(d)$	$O(d)$
	Sequences	nm^2	$O(nm^2)$	$O(nm^2)$	N/A
Arborescence	Spanning trees	$n(n-1)$	$O(n^2)$	$O(n^3)$	N/A
	Alignments	nm	$O(nm)$	$O(nm)$	N/A
Permutahedron	Permutations	d	$O(d \log d)$	N/A	$O(d \log d)$
Birkhoff	Permutation matrices	n^2	$O(n^3)$	#P-complete	$O(n^2/\epsilon)$

can similarly be organized as a $n \times m \times m$ real tensor, such that $\theta_{t,i,j} = \phi_t(\mathbf{x}_t, i, j)$, where ϕ_t is a potential function. Using the above binary tensor representation, the Frobenius inner product $\langle \boldsymbol{\theta}, \mathbf{y} \rangle$ is equal to $\sum_{t=1}^n \phi_t(\mathbf{x}_t, s_t, s_{t-1})$, the cumulated score of \mathbf{s} .

MAP inference, $\operatorname{argmax}_{\mathbf{y} \in \mathcal{Y}} \langle \mathbf{y}, \boldsymbol{\theta} \rangle$, seeks the highest-scoring sequence and can be computed using Viterbi’s algorithm (Viterbi, 1967) in $O(nm^2)$ time. Marginal inference can be computed in the same cost using the forward-backward algorithm (Baum and Petrie, 1966) or equivalently using backpropagation (Eisner, 2016; Mensch and Blondel, 2018).

When Ω is defined over $\operatorname{conv}(\mathcal{Y})$ as in (26), the Fenchel-Young loss corresponds to a linear-chain conditional random field loss (Lafferty et al., 2001). In recent years, this loss has been used to train various end-to-end natural language pipelines based on neural networks (Collobert et al., 2011; Lample et al., 2016).

Alignments. Let $\mathbf{A} \in \mathbb{R}^{m \times p}$ and $\mathbf{B} \in \mathbb{R}^{n \times p}$ be two time-series of lengths m and n , respectively. We denote by $\mathbf{a}_i \in \mathbb{R}^p$ and $\mathbf{b}_j \in \mathbb{R}^p$ their i^{th} and j^{th} observations. Our goal is to find an alignment between \mathbf{A} and \mathbf{B} , matching their observations. We define $\boldsymbol{\theta}$ as a $m \times n$ matrix, such that $\theta_{i,j}$ is the similarity between observations \mathbf{a}_i and \mathbf{b}_j . Likewise, we represent an alignment \mathbf{y} as a $m \times n$ binary matrix, such that $y_{i,j} = 1$ if \mathbf{a}_i is aligned with \mathbf{b}_j , and 0 otherwise. We write \mathcal{Y} the set of all monotonic alignment matrices, such that the path that connects the upper-left $(1, 1)$ matrix entry to the lower-right (m, n) one uses only $\downarrow, \rightarrow, \searrow$ moves.

In this setting, MAP inference, $\operatorname{argmax}_{\mathbf{y} \in \mathcal{Y}} \langle \mathbf{y}, \boldsymbol{\theta} \rangle$, corresponds to seeking the maximal similarity (or equivalently, minimal cost) alignment between the two time-series. It can be computed in $O(mn)$ time using dynamic time warping, DTW (Sakoe and Chiba, 1978). Marginal inference can be computed in the same cost by backpropagation, if we replace the hard minimum with a soft one in the DTW recursion. This algorithm is known as soft-DTW (Cuturi and Blondel, 2017; Mensch and Blondel, 2018).

Structured SVMs were combined with DTW to learn to predict music-to-score alignments (Garreau et al., 2014). Our framework easily enables extensions of this work, such as replacing DTW with soft-DTW, which amounts to introducing entropic regularization w.r.t. the probability distribution over alignments.

Spanning trees. When \mathcal{Y} is the set of possible directed spanning trees (arborescences) of a complete graph \mathcal{G} with n vertices, the convex hull $\text{conv}(\mathcal{Y}) \subset \mathbb{R}^{n(n-1)}$ is known as the arborescence polytope (Martins et al., 2009b) (each $\mathbf{y} \in \mathcal{Y}$ is a binary vector which indicates which arcs belong to the arborescence). MAP inference may be performed by maximal arborescence algorithms (Chu and Liu, 1965; Edmonds, 1967) in $O(n^2)$ time (Tarjan, 1977), and the Matrix-Tree theorem (Kirchhoff, 1847) provides an $O(n^3)$ marginal inference algorithm (Koo et al., 2007; Smith and Smith, 2007). Spanning tree structures are often used in natural language processing for (non-projective) *dependency parsing*, with graph edges corresponding to dependencies between the words in a sentence. (McDonald et al., 2005; Kiperwasser and Goldberg, 2016).

Permutations. We view ranking as a structured prediction problem. Let \mathcal{Y} be the set of d -permutations of a prescribed vector $\mathbf{w} \in \mathbb{R}^d$, i.e., $\mathcal{Y} = \{[w_{\pi_1}, \dots, w_{\pi_d}] \in \mathbb{R}^d : \pi \in \Pi\}$, where Π denotes the permutations of $(1, \dots, d)$. We assume without loss of generality that \mathbf{w} is sorted in descending order, i.e., $w_1 \geq \dots \geq w_d$. MAP inference seeks the permutation of \mathbf{w} whose inner product with $\boldsymbol{\theta}$ is maximized:

$$\max_{\mathbf{y} \in \mathcal{Y}} \langle \boldsymbol{\theta}, \mathbf{y} \rangle = \max_{\mathbf{y} \in \mathcal{Y}} \sum_{i=1}^d \theta_i y_i = \max_{\pi \in \Pi} \sum_{i=1}^d \theta_i w_{\pi_i} = \max_{\pi \in \Pi} \sum_{i=1}^d \theta_{\pi_i} w_i.$$

Since \mathbf{w} is assumed sorted, an optimal solution π^* of the last optimization problem is a permutation sorting $\boldsymbol{\theta}$ in descending order. The function $\boldsymbol{\theta} \mapsto \sum_{i=1}^d \theta_{\pi_i^*} w_i$ is called ordered weighted averaging (OWA) operator (Yager, 1988) and includes the mean and max operator as special cases. MAP inference over \mathcal{Y} can be seen as the variational formulation of the OWA operator. An optimal solution \mathbf{y}^* is simply \mathbf{w} sorted using the inverse permutation of π^* . The overall computational cost is therefore $O(d \log d)$, for sorting $\boldsymbol{\theta}$.

The convex hull of \mathcal{Y} , $\text{conv}(\mathcal{Y})$, is known as the *permutahedron* when $\mathbf{w} = [d, \dots, 1]$. For arbitrary \mathbf{w} , we follow Lim and Wright (2016) and call $\text{conv}(\mathcal{Y})$ the permutahedron induced by \mathbf{w} . Its vertices correspond to permutations of \mathbf{w} .

We can define mean-regularized ranking prediction functions $\hat{\mathbf{y}}_{\Omega}(\boldsymbol{\theta})$ if we choose \mathbf{w} such that each w_i represents the “preference” of being in position i . Intuitively, the score vector $\boldsymbol{\theta} \in \mathbb{R}^d$ should be such that each θ_i is the score of instance i (e.g., a document or a label). We give several examples below.

- Choosing $\mathbf{w} = [1, 0, \dots, 0]$, $\text{conv}(\mathcal{Y})$ is equal to Δ^d , the probability simplex. We thus recover probabilistic classification as a natural special case.
- Choosing $\mathbf{w} = \frac{1}{k} \underbrace{[1, \dots, 1]}_{k \text{ times}}, 0, \dots, 0]$, $\text{conv}(\mathcal{Y})$ is equal to $\{\mathbf{p} \in \Delta^d : \|\mathbf{p}\|_{\infty} \leq \frac{1}{k}\}$, sometimes referred to as the capped probability simplex (Warmuth and Kuzmin, 2008; Lim and Wright, 2016). This setting corresponds to predicting k -subsets.




Figure 6: **Examples of instances of permutahedron induced by w .** Round circles indicate vertices of the permutahedron, permutations of w . Choosing $w = [1, 0, 0]$ recovers the probability simplex (red) while choosing $w = [\frac{1}{2}, \frac{1}{2}, 0]$ recovers the capped probability simplex (blue). The other two instances are obtained by $w = [\frac{2}{3}, \frac{1}{3}, 0]$ (gray) and $w = [\frac{1}{2}, \frac{1}{3}, \frac{1}{6}]$ (green). Euclidean projection onto these polytopes can be cast as isotonic regression. More generally, Bregman projection reduces to isotonic *optimization*.

- Choosing $w = \frac{2}{d(d+1)}[d, d-1, \dots, 1]$ corresponds to predicting full rankings.
- Choosing $w = \frac{2}{k(k+1)}[k, k-1, \dots, 1, \underbrace{0, \dots, 0}_{d-k-1 \text{ times}}]$ corresponds to predicting partial rankings.

The corresponding polytopes are illustrated in Figure 6. In all the examples above, $w \in \Delta^d$, implying $\text{conv}(\mathcal{Y}) \subseteq \Delta^d$. Therefore, $\hat{y}_\Omega(\theta)$ outputs a probability distribution.

As discussed in §3.2, computing the regularized prediction function $\hat{y}_\Omega(\theta)$ is equivalent to a Bregman projection when $\Omega = \Psi + I_{\mathcal{C}}$, where Ψ is Legendre type. The Euclidean projection onto $\mathcal{C} = \text{conv}(\mathcal{Y})$ reduces to isotonic regression (Zeng and Figueiredo, 2015; Negrinho and Martins, 2014). The computational cost is $O(d \log d)$. More generally, Bregman projections reduce to isotonic *optimization* (Lim and Wright, 2016). This provides a unified way to compute $\hat{y}_\Omega(\theta)$ efficiently, regardless of w .

The generated Fenchel-Young loss $L_\Omega(\theta; \mathbf{y})$, is illustrated in Figure 7 for various choices of Ω . When $\Omega = 0$, as expected, the loss is zero as long as the predicted ranking is correct. Note that in order to define a meaningful loss, it is necessary that $\mathbf{y} \in \mathcal{Y}$ or more generally $\mathbf{y} \in \text{conv}(\mathcal{Y})$. That is, \mathbf{y} should belong to the convex hull of the permutations of w .

Permutahedra have been used to derive online learning to rank algorithms (Yasutake et al., 2011; Ailon et al., 2016) but it is not obvious how to extract a loss from these works. Ordered weighted averaging (OWA) operators have been used to define related top- k multiclass losses (Usunier et al., 2009; Lapin et al., 2015) but without identifying the connection




Figure 7: **Ranking losses** generated by restricting $\text{dom}(\Omega)$ to the permutahedron induced by $\mathbf{w} = [3, 2, 1]$. We set the ground-truth $\mathbf{y} = [2, 1, 3]$, i.e., a permutation of \mathbf{w} . We set $\boldsymbol{\theta} = \mathbf{y}$ and inspect how the loss changes when varying each θ_i . When $\Omega = 0$, the loss is zero as expected when $\theta_1 \in [1, 3]$, $\theta_2 \in (-\infty, 2]$ and $\theta_3 \in [2, \infty)$, since the ground-truth ranking is still correctly predicted in these intervals. When $\Omega = \frac{1}{2} \|\cdot\|^2$ and $\Omega = -H^S$, we obtain a smooth approximation. Although not done here for clarity, if \mathbf{w} is normalized such that $\mathbf{w} \in \Delta^d$, then $\hat{\mathbf{y}}_{\Omega}(\boldsymbol{\theta}) \in \Delta^d$ as well.

with permutahedra. Our construction follows directly from the general Fenchel-Young loss framework and provides a novel geometric perspective.

Permutation matrices. Let $\mathcal{Y} \subset \{0, 1\}^{n \times n}$ be the set of $n \times n$ permutation matrices. MAP inference is the solution of the linear assignment problem

$$\max_{\mathbf{y} \in \mathcal{Y}} \langle \boldsymbol{\theta}, \mathbf{y} \rangle = \max_{\boldsymbol{\pi} \in \Pi} \sum_{i=1}^n \theta_{i, \pi_i},$$

where Π denotes the permutations of $(1, \dots, n)$. The problem can be solved exactly in $O(n^3)$ time using the Hungarian algorithm (Kuhn, 1955) or the Jonker-Volgenant algorithm (Jonker and Volgenant, 1987). Noticeably, marginal inference is known to be $\#\text{P}$ -complete (Valiant, 1979; Taskar, 2004, Section 3.5). This makes it an open problem how to solve marginal inference for this polytope.

In contrast, projections on the convex hull of \mathcal{Y} , the set of doubly stochastic matrices known as Birkhoff polytope (Birkhoff, 1946), can be computed efficiently. Since the Birkhoff polytope is a special case of transportation polytope, we can leverage algorithms from the optimal transport literature to compute the mean-regularized prediction function. When $\Omega(\boldsymbol{\mu}) = \langle \boldsymbol{\mu}, \log \boldsymbol{\mu} \rangle$, $\hat{\mathbf{y}}_{\Omega}(\boldsymbol{\theta})$ can be computed using the Sinkhorn algorithm (Sinkhorn and Knopp, 1967; Cuturi, 2013). For other regularizers, we can use Dykstra’s algorithm (Dessein et al., 2016) or dual approaches (Blondel et al., 2018). The cost of obtaining an ϵ -approximate solution is typically $O(n^2/\epsilon)$.

The Birkhoff polytope has been used to define continuous relaxations of non-convex ranking losses (Adams and Zemel, 2011). In contrast, the Fenchel-Young loss over the Birkhoff polytope (new to our knowledge) is convex by construction. Although working with permutation matrices is more computationally expensive than working with permutations, it brings different modeling power, since it allows to take into account similarity between instances

(e.g., text documents) through the similarity matrix θ . It also enables other applications, such as learning to match.

7.4 Structured separation margins

We end this section by extending some of the results in §5 for the structured prediction case, showing that **there is also a relation between structured entropies and margins**. In the sequel, we assume that structured outputs are contained in a sphere of radius r , i.e., $\|\mathbf{y}\| = r$ for any $\mathbf{y} \in \mathcal{Y}$. This holds for all the examples above (sequences, alignments, spanning trees, permutations of a given vector, and permutation matrices). In particular, it holds whenever outputs are represented as binary vectors with a constant number of entries set to 1; this includes overcomplete parametrizations of discrete graphical models (Wainwright and Jordan, 2008).

Definition 5 *Structured separation margin*

Let $L(\theta; \mathbf{y})$ be a loss function over $\mathbb{R}^d \times \mathcal{Y}$. We say that L has the structured separation margin property if there exists $m > 0$ such that, for any $\mathbf{y} \in \mathcal{Y}$:

$$\langle \theta, \mathbf{y} \rangle \geq \max_{\mathbf{y}' \in \mathcal{Y}} \left(\langle \theta, \mathbf{y}' \rangle + \frac{m}{2} \|\mathbf{y} - \mathbf{y}'\|^2 \right) \Rightarrow L(\theta; \mathbf{y}) = 0. \quad (27)$$

The smallest possible m that satisfies (27) is called the margin of L , denoted $\text{margin}(L)$.

Note that this definition generalizes the unstructured case (Definition 4), which is recovered when $\mathcal{Y} = \{\mathbf{e}_i\}_{i=1}^d$. Note also that, when outputs are represented as binary vectors, the term $\|\mathbf{y} - \mathbf{y}'\|^2$ is a **Hamming distance**, which counts how many bits need to be flipped to transform \mathbf{y}' into \mathbf{y} . The most famous example of a loss with a structured separation margin in the **structured hinge loss** used in structured support vector machines (Taskar, 2004; Tsochantaridis et al., 2005).

The next proposition extends Proposition 5. We defer its proof to §B.7.

Proposition 8 *Assume Ω is convex and \mathcal{Y} is contained in a sphere of radius r . Then:*

1. The loss L_Ω has a structured separation margin iff there is a $m > 0$ such that, for any $\mathbf{y} \in \mathcal{Y}$, $m\mathbf{y} \in \partial\Omega(\mathbf{y})$.
2. If the above holds, then the margin of L_Ω is given by the smallest such m or, equivalently,

$$\text{margin}(L_\Omega) = \sup_{\mu \in \text{conv}(\mathcal{Y}), \mathbf{y} \in \mathcal{Y}} \frac{\Omega(\mathbf{y}) - \Omega(\mu)}{r^2 - \langle \mathbf{y}, \mu \rangle}. \quad (28)$$

Unit margin of the SparseMAP loss. We can invoke Proposition 8 to show that the SparseMAP loss of Niculae et al. (2018) has a structured margin of 1, a novel result that follows directly from our construction. Indeed, from (28), we have that, for $\Omega(\mu) = \frac{1}{2}\|\mu\|^2$:

$$\text{margin}(L_\Omega) = \sup_{\mu \in \text{conv}(\mathcal{Y}), \mathbf{y} \in \mathcal{Y}} \frac{\frac{1}{2}r^2 - \frac{1}{2}\|\mu\|^2}{r^2 - \langle \mathbf{y}, \mu \rangle} = 1 - \inf_{\mu \in \text{conv}(\mathcal{Y}), \mathbf{y} \in \mathcal{Y}} \frac{\frac{1}{2}\|\mathbf{y} - \mu\|^2}{\langle \mathbf{y}, \mathbf{y} - \mu \rangle} \leq 1,$$




Figure 8: **Summary of computational oracles:** $\hat{\mathbf{y}}(\boldsymbol{\theta})$ is used to compute unregularized predictions or as a linear maximization oracle in conditional gradient algorithms; $\hat{\mathbf{y}}_{\Omega}(\boldsymbol{\theta})$ is used to compute regularized predictions, loss values $L_{\Omega}(\boldsymbol{\theta}; \mathbf{y}) = f_{\boldsymbol{\theta}}(\mathbf{y}) - f_{\boldsymbol{\theta}}(\hat{\mathbf{y}}_{\Omega}(\boldsymbol{\theta}))$, where $f_{\boldsymbol{\theta}}(\boldsymbol{\mu}) := \Omega(\boldsymbol{\mu}) - \langle \boldsymbol{\theta}, \boldsymbol{\mu} \rangle$, and loss gradients $\nabla L_{\Omega}(\boldsymbol{\theta}; \mathbf{y}) = \hat{\mathbf{y}}_{\Omega}(\boldsymbol{\theta}) - \mathbf{y}$; finally, $\operatorname{prox}_{\Omega}(\boldsymbol{\theta})$ is used by training algorithms, in particular block coordinate ascent algorithms. In the above, we assume Ψ is Legendre-type and $\mathcal{C} \subseteq \operatorname{dom}(\Psi)$.

where the last inequality follows from the fact that both the numerator and denominator in the second term are non-negative, the latter due to the Cauchy-Schwartz inequality. We now show that, for any $\mathbf{y} \in \mathcal{Y}$, we have $\inf_{\boldsymbol{\mu} \in \operatorname{conv}(\mathcal{Y})} \frac{\frac{1}{2} \|\mathbf{y} - \boldsymbol{\mu}\|^2}{\langle \mathbf{y}, \mathbf{y} - \boldsymbol{\mu} \rangle} = 0$. Choosing $\boldsymbol{\mu} = t\mathbf{y}' + (1-t)\mathbf{y}$, for an arbitrary $\mathbf{y}' \in \mathcal{Y} \setminus \{\mathbf{y}\}$, and letting $t \rightarrow 0^+$, we obtain $\frac{\frac{1}{2} \|\mathbf{y} - \boldsymbol{\mu}\|^2}{\langle \mathbf{y}, \mathbf{y} - \boldsymbol{\mu} \rangle} = \frac{\frac{1}{2} t \|\mathbf{y} - \mathbf{y}'\|^2}{\langle \mathbf{y}, \mathbf{y} - \mathbf{y}' \rangle} \rightarrow 0$.

8. Algorithms for learning with Fenchel-Young losses

In this section, we present generic primal (§8.1) and dual (§8.2) algorithms for training predictive models with a Fenchel-Young loss L_{Ω} for arbitrary Ω . In doing so, we obtain unified algorithms for training models with a wealth of existing and new loss functions. We then discuss algorithms for computing regularized prediction functions (§8.3) and proximity operators (§8.4). We summarize these “computational oracles” in Figure 8.

8.1 Primal training

Let $\mathbf{f}_W : \mathcal{X} \rightarrow \mathbb{R}^d$ be a model parametrized by W . To learn W from training data $\{(\mathbf{x}_i, \mathbf{y}_i)\}_{i=1}^n$, $(\mathbf{x}_i, \mathbf{y}_i) \in \mathcal{X} \times \mathcal{Y}$, we minimize the regularized empirical risk

$$\min_W F(W) + G(W) := \sum_{i=1}^n L_{\Omega}(\mathbf{f}_W(\mathbf{x}_i); \mathbf{y}_i) + G(W), \quad (29)$$

where G is a regularization function w.r.t. parameters W . Typical examples are Tikhonov regularization, $G(W) = \frac{\lambda}{2} \|W\|_F^2$, and elastic net regularization, $G(W) = \frac{\lambda}{2} \|W\|_F^2 + \lambda \rho \|W\|$, for some $\lambda > 0$ and $\rho \geq 0$.

The objective in (29) is very broad and allows to learn, e.g., linear models or neural networks using Fenchel-Young losses. Since L_{Ω} is convex, if \mathbf{f}_W is a linear model and G is convex,

then (29) is convex as well. Assuming further that Ω is strongly convex, then L_Ω is smooth and (29) can be solved globally using proximal gradient algorithms (Wright et al., 2009; Beck and Teboulle, 2009; Defazio et al., 2014). From Proposition 2, the gradient of $L_\Omega(\boldsymbol{\theta}; \mathbf{y})$ w.r.t. $\boldsymbol{\theta}$ is the residual vector

$$\nabla L_\Omega(\boldsymbol{\theta}; \mathbf{y}) = \widehat{\mathbf{y}}_\Omega(\boldsymbol{\theta}) - \mathbf{y} \in \mathbb{R}^d.$$

Using the chain rule, the gradient of F w.r.t. W is

$$\nabla F(W) = \sum_{i=1}^n D_{\mathbf{f}_W(\mathbf{x}_i)}^\top \nabla L_\Omega(\mathbf{f}_W(\mathbf{x}_i); \mathbf{y}_i),$$

where $D_{\mathbf{f}_W(\mathbf{x}_i)}$ is the Jacobian of $\mathbf{f}_W(\mathbf{x}_i)$ w.r.t. W , a linear map from the space of W to \mathbb{R}^d . For linear models, we set $\boldsymbol{\theta}_i = \mathbf{f}_W(\mathbf{x}_i) = W\mathbf{x}_i$, where $W \in \mathbb{R}^{d \times p}$ and $\mathbf{x}_i \in \mathbb{R}^p$, and thus get

$$\nabla F(W) = (\widehat{Y}_\Omega - Y)^\top X,$$

where \widehat{Y}_Ω , Y and X are matrices whose rows gather $\widehat{\mathbf{y}}_\Omega(W\mathbf{x}_i)$, \mathbf{y}_i and \mathbf{x}_i , for $i = 1, \dots, n$.

In summary, the two main computational ingredients to solve (29) are $\widehat{\mathbf{y}}_\Omega$ and the proximity operator of G (when G is non-differentiable). This separation of concerns results in very modular implementations.

Remark on temperature scaling. When \mathbf{f}_W is a linear model and G is a homogeneous function, it is easy to check that the temperature scaling parameter $t > 0$ from Proposition 2 and the regularization strength $\lambda > 0$ above are redundant. Hence, we can set $t = 1$ at training time, without loss of generality. However, at test time, using $\widehat{\mathbf{y}}_{t\Omega}(\boldsymbol{\theta})$ and tuning t can potentially improve accuracy. For sparse prediction functions, t also gives control over sparsity (the smaller, the sparser).

8.2 Dual training

We now derive dual training of Fenchel-Young losses. Although our treatment follows closely Shalev-Shwartz and Zhang (2016), it is different in that we put the output regularization Ω at the center of all computations, leading to a new perspective.

Dual objective. For linear models $\mathbf{f}_W(\mathbf{x}) = W\mathbf{x}$, where $W \in \mathbb{R}^{d \times p}$, it can sometimes be more computationally efficient to solve the corresponding dual problem, especially in the $n \ll p$ setting. From Fenchel's duality theorem (see for instance Borwein and Lewis (2010, Theorem 3.3.5)), we find that the dual problem of (29) is

$$\max_{\alpha \in \mathbb{R}^{n \times d}} - \sum_{i=1}^n L_\Omega^*(-\alpha_i; \mathbf{y}_i) - G^*(\alpha^\top X),$$

and where L_Ω^* is the convex conjugate of L_Ω in its first argument.

We now rewrite the dual problem using the specific form of L_Ω . Using $L_\Omega^*(-\boldsymbol{\alpha}; \mathbf{y}) = \Omega(\mathbf{y} - \boldsymbol{\alpha}) - \Omega(\mathbf{y})$ and using the change of variable $\boldsymbol{\mu}_i := \mathbf{y}_i - \boldsymbol{\alpha}_i$, we get

$$\max_{\boldsymbol{\mu} \in \mathbb{R}^{n \times d}} -D(\boldsymbol{\mu}) \text{ s.t. } \boldsymbol{\mu}_i \in \text{dom}(\Omega) \forall i \in [n], \quad (30)$$

where we defined

$$D(\boldsymbol{\mu}) := \sum_{i=1}^n \Omega(\boldsymbol{\mu}_i) - \Omega(\mathbf{y}_i) + G^*(V(\boldsymbol{\mu})) \quad \text{and} \quad V(\boldsymbol{\mu}) := (Y - \boldsymbol{\mu})^\top X.$$

This expression is informative as we can interpret each $\boldsymbol{\mu}_i$ as regularized predictions, i.e., $\boldsymbol{\mu}_i$ should belong to $\text{dom}(\Omega)$, the same domain as $\widehat{\boldsymbol{y}}_\Omega$. The fact that $G^*(V(\boldsymbol{\mu}))$ is a function of the predictions $\boldsymbol{\mu}$ is similar to the value regularization framework of Rifkin and Lippert (2007). A key difference with the regularization Ω , however, is that $G^*(V(\boldsymbol{\mu}))$ depends on the training data X through $V(\boldsymbol{\mu})$. When $G(W) = \frac{\lambda}{2} \|W\|_F^2 \Leftrightarrow G^*(V) = \frac{1}{2\lambda} \|V\|_F^2$, we obtain $G^*(V(\boldsymbol{\mu})) = \frac{1}{2\lambda} \text{Trace}((Y - \boldsymbol{\mu})^\top K(Y - \boldsymbol{\mu}))$, where $K := XX^\top$ is the Gram matrix.

Primal-dual relationship. Assuming G is a λ -strongly convex regularizer, given an optimal dual solution $\boldsymbol{\mu}^*$ to (30), we may retrieve the optimal primal solution W^* by

$$W^* = \nabla G^*(V(\boldsymbol{\mu}^*)).$$

Coordinate ascent. We can solve (30) using block coordinate ascent algorithms. At every iteration, we pick $i \in [n]$ and consider the following sub-problem associated with i :

$$\underset{\boldsymbol{\mu}_i \in \text{dom}(\Omega)}{\text{argmin}} \quad \Omega(\boldsymbol{\mu}_i) + G^*(V(\boldsymbol{\mu})). \quad (31)$$

Since this problem could be hard to solve, we follow Shalev-Shwartz and Zhang (2016, Option I) and consider instead a quadratic upper-bound. If G is λ -strongly convex, G^* is $\frac{1}{\lambda}$ -smooth w.r.t. the dual norm $\|\cdot\|$ and it holds that

$$G^*(V(\boldsymbol{\mu})) \leq G^*(V(\bar{\boldsymbol{\mu}})) + \langle \nabla G^*(V(\bar{\boldsymbol{\mu}})), V(\boldsymbol{\mu}) - V(\bar{\boldsymbol{\mu}}) \rangle + \frac{1}{2\lambda} \|V(\boldsymbol{\mu}) - V(\bar{\boldsymbol{\mu}})\|^2,$$

where $\bar{\boldsymbol{\mu}}$ denotes the current iterate of $\boldsymbol{\mu}$. Using $V(\boldsymbol{\mu}) - V(\bar{\boldsymbol{\mu}}) = \sum_{i=1}^n (\bar{\boldsymbol{\mu}}_i - \boldsymbol{\mu}_i) \mathbf{x}_i^\top$, we get

$$G^*(V(\boldsymbol{\mu})) \leq G^*(V(\bar{\boldsymbol{\mu}})) + \sum_{i=1}^n \langle \nabla G^*(V(\bar{\boldsymbol{\mu}})) \mathbf{x}_i, \bar{\boldsymbol{\mu}}_i - \boldsymbol{\mu}_i \rangle + \frac{\sigma_i}{2} \|\bar{\boldsymbol{\mu}}_i - \boldsymbol{\mu}_i\|^2,$$

where $\sigma_i := \frac{\|\mathbf{x}_i\|^2}{\lambda}$. Substituting $G^*(V(\boldsymbol{\mu}))$ by the above upper bound into (31) and ignoring constant terms, we get the following approximate sub-problem:

$$\underset{\boldsymbol{\mu}_i \in \text{dom}(\Omega)}{\text{argmin}} \quad \Omega(\boldsymbol{\mu}_i) - \boldsymbol{\mu}_i^\top \mathbf{v}_i + \frac{\sigma_i}{2} \|\boldsymbol{\mu}_i\|^2 = \text{prox}_{\frac{1}{\sigma_i} \Omega} \left(\frac{\mathbf{v}_i}{\sigma_i} \right), \quad (32)$$

where $\mathbf{v}_i := \nabla G^*(V(\bar{\boldsymbol{\mu}})) \mathbf{x}_i + \sigma_i \bar{\boldsymbol{\mu}}_i$. Note that when G^* is a quadratic function, (32) is an optimal solution of (31).

Examples of parameter regularization G . We now give examples of $G^*(W)$ and $\nabla G^*(W)$ for two commonly used regularization: squared ℓ_2 and elastic-net regularization.

When $G(W) = \frac{\lambda}{2}\|W\|_F^2$, we obtain

$$G^*(V) = \frac{1}{2\lambda}\|V\|_F^2 \quad \text{and} \quad \nabla G^*(V) = \frac{1}{\lambda}V.$$

When $G(W) = \frac{\lambda}{2}\|W\|_F^2 + \lambda\rho R(W)$, for some other regularizer $R(W)$, we obtain

$$\begin{aligned} \nabla G^*(V) &= \operatorname{argmax}_W \langle W, V \rangle - \frac{\lambda}{2}\|W\|_F^2 - \lambda\rho R(W) \\ &= \operatorname{argmin}_W \frac{1}{2}\left\|W - \frac{V}{\lambda}\right\|_F^2 + \rho R(W) \\ &= \operatorname{prox}_{\rho R}(V/\lambda). \end{aligned}$$

The conjugate is equal to $G^*(V) = \langle \nabla G^*(V), V \rangle - G(\nabla G^*(V))$.

For instance, if we choose $R(W) = \|W\|_1$, then $G(W)$ is the **elastic-net** regularization and $\operatorname{prox}_{\rho R}$ is the well-known **soft-thresholding** operator

$$\operatorname{prox}_{\rho R}(V) = \operatorname{sign}(V)[|V| - \rho]_+,$$

where all operations above are performed element-wise.

Summary. To summarize, on each iteration, we pick $i \in [n]$ and perform the update

$$\boldsymbol{\mu}_i \leftarrow \operatorname{prox}_{\frac{1}{\sigma_i}\Omega}(\mathbf{v}_i/\sigma_i),$$

where we defined $\mathbf{v}_i := \nabla G^*(V(\boldsymbol{\mu}))\mathbf{x}_i + \sigma_i\boldsymbol{\mu}_i$ and $\sigma_i := \frac{\|\mathbf{x}_i\|^2}{\lambda}$. This update does not require choosing any learning rate. Interestingly, if $\operatorname{prox}_{\tau\Omega}$ is sparse, then **so are the dual variables** $\{\boldsymbol{\mu}_i\}_{i=1}^n$. Sparsity in the dual variables is particularly useful when kernelizing models, as it makes computing predictions more efficient.

When i is picked uniformly at random, this block coordinate ascent algorithm is known to converge to an optimal solution W^* , at a linear rate if L_Ω is smooth (i.e., if Ω is strongly convex) (Shalev-Shwartz and Zhang, 2016). When $\operatorname{dom}(\Omega)$ is a compact set, another option to solve (30) that does not involve proximity operators is the block Frank-Wolfe algorithm (Lacoste-Julien et al., 2012).

8.3 Regularized prediction functions

The regularized prediction function $\widehat{\mathbf{y}}_\Omega(\boldsymbol{\theta})$ does not generally enjoy a closed-form expression and one must resort to generic algorithms to compute it. In this section, we first discuss two such algorithms: projected gradient and conditional gradient methods. Then, we present a new more efficient algorithm when $\operatorname{dom}(\Omega) \subseteq \Delta^d$ and Ω is uniformly separable.

Generic algorithms. In their greater generality, regularized prediction functions involve the following optimization problem

$$\min_{\boldsymbol{\mu} \in \text{dom}(\Omega)} f_{\boldsymbol{\theta}}(\boldsymbol{\mu}) := \Omega(\boldsymbol{\mu}) - \langle \boldsymbol{\theta}, \boldsymbol{\mu} \rangle. \quad (33)$$

Assuming Ω is convex and smooth (differentiable and with Lipschitz-continuous gradient), we can solve this problem to arbitrary precision using projected gradient methods. Unfortunately, the projection onto $\text{dom}(\Omega)$, which is needed by these algorithms, is often as challenging to solve as (33) itself. This is especially the case when $\text{dom}(\Omega)$ is $\text{conv}(\mathcal{Y})$, the convex hull of a combinatorial set of structured objects.

Assuming further that $\text{dom}(\Omega)$ is a compact set, an appealing alternative which sidesteps these issues is provided by conditional gradient (a.k.a. Frank-Wolfe) algorithms (Frank and Wolfe, 1956; Dunn and Harshbarger, 1978; Jaggi, 2013). Their main advantage stems from the fact that they access $\text{dom}(\Omega)$ only through the linear maximization oracle $\text{argmax}_{\mathbf{y} \in \text{dom}(\Omega)} \langle \boldsymbol{\theta}, \mathbf{y} \rangle$. CG algorithms maintain the current solution as a sparse convex combination of vertices of $\text{dom}(\Omega)$. At each iteration, the linear maximization oracle is used to pick a new vertex to add to the combination. Despite its simplicity, the procedure converges to an optimal solution, albeit at a sub-linear rate (Jaggi, 2013). Linear convergence rates can be obtained using away-step, pairwise and full-corrective variants (Lacoste-Julien and Jaggi, 2015). When Ω is a quadratic function (as in the case of SparseMAP), an approximate correction step can achieve finite convergence efficiently. The resulting algorithm is known as the active set method (Nocedal and Wright, 1999, chapters 16.4 & 16.5), and is a generalization of Wolfe’s min-norm point algorithm (Wolfe, 1976). See also Vinyes and Obozinski (2017) for a more detailed discussion on these algorithms and Niculae et al. (2018) for an instantiation, in the specific case of SparseMAP.

Although not explored in this paper, similar algorithms have been developed to minimize a function over conic hulls (Locatello et al., 2017). Such algorithms allow to compute a regularized prediction function that outputs a conic combination of elementary structures, instead of a convex one. It can be seen as the structured counterpart of the regularized prediction function over positive measures (§6).

Reduction to root finding. When Ω is a strictly convex regularization function over Δ^d and is uniformly separable, i.e., $\Omega(\mathbf{p}) = \sum_{i=1}^d g(p_i)$ for some strictly convex function g , we now show that $\hat{\mathbf{y}}_{\Omega}(\boldsymbol{\theta})$ can be computed in linear time.

Proposition 9 *Reduction to root finding*

Let $g: [0, 1] \rightarrow \mathbb{R}_+$ be a strictly convex and differentiable function. Then,

$$\hat{\mathbf{y}}_{\Omega}(\boldsymbol{\theta}) = \text{argmax}_{\mathbf{p} \in \Delta^d} \langle \mathbf{p}, \boldsymbol{\theta} \rangle - \sum_{j=1}^d g(p_j) = \mathbf{p}(\tau^*)$$

where

$$\mathbf{p}(\tau) := (g')^{-1}(\max\{\boldsymbol{\theta} - \tau, g'(0)\})$$

and where τ^* is a root of $\phi(\tau) := \langle \mathbf{p}(\tau), \mathbf{1} \rangle - 1$.

Moreover, τ^* belongs to the tight search interval $[\tau_{\min}, \tau_{\max}]$, where

$$\tau_{\min} := \max(\boldsymbol{\theta}) - g'(1) \quad \text{and} \quad \tau_{\max} := \max(\boldsymbol{\theta}) - g'(1/d).$$

The equation of $\mathbf{p}(\tau)$ can be seen as a generalization of (7). An approximate τ such that $|\phi(\tau)| \leq \epsilon$ can be found in $O(1/\log \epsilon)$ time by, e.g., bisection. The related problem of Bregman projection onto the probability simplex was recently studied by Krichene et al. (2015) but our derivation is different and more direct (cf. §B.8).

For example, when Ω is the negative α -Tsallis entropy $-\mathbf{H}_\alpha^T$, which we saw can be written in separable form in (18), we obtain

$$g(t) = \frac{t^\alpha - t}{\alpha(\alpha - 1)}, \quad g'(t) = \frac{t^{\alpha-1} - 1}{\alpha - 1} \quad \text{and} \quad (g')^{-1}(s) = (1 + (\alpha - 1)s)^{\frac{1}{\alpha-1}},$$

yielding

$$\mathbf{p}(\tau) = \left(1 + (\alpha - 1) \max(\boldsymbol{\theta} - \tau, -1/\alpha-1)\right)^{\frac{1}{\alpha-1}}.$$

From the root τ^* of $\phi(\tau) = \langle \mathbf{p}(\tau), \mathbf{1} \rangle - 1$, we obtain $\hat{\mathbf{y}}_{-\mathbf{H}_\alpha^T}(\boldsymbol{\theta}) = \mathbf{p}(\tau^*)$.

8.4 Proximity operators

Computing the proximity operator $\text{prox}_{\tau\Omega}(\boldsymbol{\eta})$, defined in (2), usually involves a more challenging optimization problem than $\hat{\mathbf{y}}_\Omega(\boldsymbol{\theta})$. For instance, when Ω is Shannon's negative entropy over Δ^d , $\hat{\mathbf{y}}_\Omega$ enjoys a closed-form solution (the softmax) but not prox_Ω . However, we can always compute $\text{prox}_{\tau\Omega}(\boldsymbol{\eta})$ by first-order gradient methods given access to $\hat{\mathbf{y}}_\Omega$. Indeed, using the Moreau decomposition, we have

$$\text{prox}_{\tau\Omega}(\boldsymbol{\eta}) = \boldsymbol{\eta} - \text{prox}_{\frac{\Omega^*}{\tau}}(\boldsymbol{\eta}/\tau).$$

Since $\text{dom}(\Omega^*) = \mathbb{R}^d$, computing $\text{prox}_{\frac{\Omega^*}{\tau}}(\boldsymbol{\eta}/\tau)$ only involves an unconstrained optimization problem, $\text{argmin}_{\boldsymbol{\theta} \in \mathbb{R}^d} \frac{1}{2} \|\boldsymbol{\theta} - \boldsymbol{\eta}/\tau\|^2 + \frac{1}{\tau} \Omega^*(\boldsymbol{\theta})$. Since $\nabla \Omega^* = \hat{\mathbf{y}}_\Omega$, that optimization problem can easily be solved by any first-order gradient method given access to $\hat{\mathbf{y}}_\Omega$. For specific choices of Ω , $\text{prox}_{\tau\Omega}(\boldsymbol{\eta})$ can be computed directly more efficiently, as we now discuss.

Closed-form expressions. We now give examples of commonly-used loss functions for which $\text{prox}_{\tau\Omega}$ enjoys a closed-form expression.

For the squared loss, we choose $\Omega(\boldsymbol{\mu}) = \frac{1}{2} \|\boldsymbol{\mu}\|^2$. Hence:

$$\text{prox}_{\tau\Omega}(\boldsymbol{\eta}) = \text{argmin}_{\boldsymbol{\mu} \in \mathbb{R}^d} \frac{1}{2} \|\boldsymbol{\mu} - \boldsymbol{\eta}\|^2 + \frac{\tau}{2} \|\boldsymbol{\mu}\|^2 = \frac{\boldsymbol{\eta}}{\tau + 1}.$$

For the perceptron loss (Rosenblatt, 1958; Collins, 2002), we choose $\Omega = I_{\Delta^d}$. Hence:

$$\text{prox}_{\tau\Omega}(\boldsymbol{\eta}) = \text{argmin}_{\mathbf{p} \in \Delta^d} \frac{1}{2} \|\mathbf{p} - \boldsymbol{\eta}\|^2.$$

For the sparsemax loss (Martins and Astudillo, 2016), we choose $\Omega = \frac{1}{2} \|\cdot\|^2 + I_{\Delta^d}$. Hence:

$$\text{prox}_{\tau\Omega}(\boldsymbol{\eta}) = \underset{\mathbf{p} \in \Delta^d}{\text{argmin}} \frac{1}{2} \|\mathbf{p} - \boldsymbol{\eta}\|^2 + \frac{\tau}{2} \|\mathbf{p}\|^2 = \underset{\mathbf{p} \in \Delta^d}{\text{argmin}} \frac{1}{2} \left\| \mathbf{p} - \frac{\boldsymbol{\eta}}{\tau + 1} \right\|^2.$$

For the cost-sensitive multiclass hinge loss, we choose $\Omega = I_{\Delta^d} - \langle \cdot, \mathbf{c}_y \rangle$. Hence:

$$\text{prox}_{\tau\Omega}(\boldsymbol{\eta}) = \underset{\mathbf{p} \in \Delta^d}{\text{argmin}} \frac{1}{2} \|\mathbf{p} - \boldsymbol{\eta}\|^2 - \tau \langle \mathbf{p}, \mathbf{c}_y \rangle = \underset{\mathbf{p} \in \Delta^d}{\text{argmin}} \frac{1}{2} \|\mathbf{p} - (\boldsymbol{\eta} + \tau \mathbf{c}_y)\|^2.$$

Choosing $\mathbf{c}_y = \mathbf{1} - \mathbf{y}$, where $\mathbf{y} \in \{\mathbf{e}_i\}_{i=1}^d$ is the ground-truth label, gives the proximity operator for the usual multiclass hinge loss (Crammer and Singer, 2001).

For the Shannon entropy and 1.5-Tsallis entropy, we show that prox_{-H} reduces to root finding in Appendix A.1.

9. Experiments

In this section, we demonstrate one of the key features of Fenchel-Young losses: their ability to induce **sparse** probability distributions. We focus on two tasks: label proportion estimation (§9.1) and dependency parsing (§9.2).

9.1 Label proportion estimation experiments

As we saw in §4, α -Tsallis entropies generate a family of losses, with the logistic ($\alpha \rightarrow 1$) and sparsemax losses ($\alpha = 2$) as important special cases. In addition, they are twice differentiable for $\alpha \in [1, 2)$, produce sparse probability distributions for $\alpha > 1$ and are computationally efficient for any $\alpha \geq 1$, thanks to Proposition 9. In this section, we demonstrate their effectiveness on the task of label proportion estimation and compare different solvers for computing the regularized prediction function $\hat{\mathbf{y}}_{-H_\alpha^\top}$.

Experimental setup. Given an input vector $\mathbf{x} \in \mathcal{X} \subseteq \mathbb{R}^p$, where p is the number of features, our goal is to estimate a vector of label proportions $\mathbf{y} \in \Delta^d$, where d is the number of classes. If \mathbf{y} is sparse, we expect the superiority of Tsallis losses over the conventional logistic loss on this task. At training time, given a set of n pairs $(\mathbf{x}_i, \mathbf{y}_i)$, we estimate a matrix $W \in \mathbb{R}^{d \times p}$ by minimizing the convex objective

$$R(W) := \sum_{i=1}^n L_\Omega(W \mathbf{x}_i; \mathbf{y}_i) + \frac{\lambda}{2} \|W\|_F^2.$$

We optimize the loss using L-BFGS (Liu and Nocedal, 1989). From Proposition 2 and using the chain rule, we obtain the gradient expression $\nabla R(W) = (\hat{Y}_\Omega - Y)^\top X + \lambda W$, where \hat{Y}_Ω , Y and X are matrices whose rows gather $\hat{\mathbf{y}}_\Omega(W \mathbf{x}_i)$, \mathbf{y}_i and \mathbf{x}_i , for $i = 1, \dots, n$. At test time, we predict label proportions by $\mathbf{p} = \hat{\mathbf{y}}_{-H_\alpha^\top}(W \mathbf{x})$.

Table 3: Dataset statistics

Dataset	Type	Train	Dev	Test	Features	Classes	Avg. labels
Birds	Audio	134	45	172	260	19	2
Cal500	Music	376	126	101	68	174	25
Emotions	Music	293	98	202	72	6	2
Mediamill	Video	22,353	7,451	12,373	120	101	5
Scene	Images	908	303	1,196	294	6	1
SIAM TMC	Text	16,139	5,380	7,077	30,438	22	2
Yeast	Micro-array	1,125	375	917	103	14	4

Real data experiments. We ran experiments on 7 standard multi-label benchmark datasets — see Table 3 for dataset characteristics¹. For all datasets, we removed samples with no label, normalized samples to have zero mean unit variance, and normalized labels to lie in the probability simplex. We chose $\lambda \in \{10^{-4}, 10^{-3}, \dots, 10^4\}$ and $\alpha \in \{1, 1.1, \dots, 2\}$ against the validation set. We report the test set mean Jensen-Shannon divergence, $\text{JS}(\mathbf{p}, \mathbf{y}) := 1/2 \text{KL}(\mathbf{p} \| \frac{\mathbf{p} + \mathbf{y}}{2}) + 1/2 \text{KL}(\mathbf{y} \| \frac{\mathbf{p} + \mathbf{y}}{2})$, and the mean squared error $1/2 \|\mathbf{p} - \mathbf{y}\|^2$ in Table 4. As can be seen, the loss with tuned α achieves the best averaged rank overall. Tuning α allows to choose the best loss in the family in a data-driven fashion.

Synthetic data experiments. We follow Martins and Astudillo (2016) and generate a document $\mathbf{x} \in \mathbb{R}^p$ from a mixture of multinomials and label proportions $\mathbf{y} \in \Delta^d$ from a multinomial. The number of words in \mathbf{x} and labels in \mathbf{y} is sampled from a Poisson distribution — see Martins and Astudillo (2016) for a precise description of the generative process. We use 1200 samples as training set, 200 samples as validation set and 1000 samples as test set. We tune $\lambda \in \{10^{-6}, 10^{-5}, \dots, 10^0\}$ and $\alpha \in \{1.0, 1.1, \dots, 2.0\}$ against the validation set. We report the Jensen-Shannon divergence in Figure 10. Results using the mean squared error (MSE) were entirely similar. When the number of classes is 10, Tsallis and sparsemax losses perform almost exactly the same, both outperforming softmax. When the number of classes is 50, Tsallis losses outperform both sparsemax and softmax.

Solver comparison. We also compared bisection (binary search) and Brent’s method for solving (5) by root finding (Proposition 9). We focus on $H_{1.5}^T$, i.e. the 1.5-Tsallis entropy, and also compare against using a generic projected gradient algorithm (FISTA) to solve (5) naively. We measure the time needed to reach a solution \mathbf{p} with $\|\mathbf{p} - \mathbf{p}^*\|_2 < 10^{-5}$, over 200 samples $\boldsymbol{\theta} \in \mathbb{R}^d \sim \mathcal{N}(\mathbf{0}, \sigma \mathbf{I})$ with $\log \sigma \sim \mathcal{U}(-4, 4)$. Median and 99% CI times reported in Figure 9 reveal that root finding scales better, with Brent’s method outperforming FISTA by one to two orders of magnitude.




Figure 9: Solver comparison.

1. The datasets can be downloaded from <http://mulan.sourceforge.net/datasets-mlc.html> and <https://www.csie.ntu.edu.tw/~cjlin/libsvmtools/datasets/>.

Table 4: **Test-set performance of Tsallis losses for various α on the task of sparse label proportion estimation:** average Jensen-Shannon divergence (left) and mean squared error (right). Lower is better.

	$\alpha = 1$ (logistic)	$\alpha = 1.5$	$\alpha = 2$ (sparsemax)	tuned α
Birds	0.359 / 0.530	0.364 / 0.504	0.364 / 0.504	0.358 / 0.501
Cal500	0.454 / 0.034	0.456 / 0.035	0.452 / 0.035	0.456 / 0.034
Emotions	0.226 / 0.327	0.225 / 0.317	0.225 / 0.317	0.224 / 0.321
Mediamill	0.375 / 0.208	0.363 / 0.193	0.356 / 0.191	0.361 / 0.193
Scene	0.175 / 0.344	0.176 / 0.363	0.176 / 0.363	0.175 / 0.345
TMC	0.225 / 0.337	0.224 / 0.327	0.224 / 0.327	0.217 / 0.328
Yeast	0.307 / 0.183	0.314 / 0.186	0.314 / 0.186	0.307 / 0.183
Avg. rank	2.57 / 2.71	2.71 / 2.14	2.14 / 2.00	1.43 / 1.86




Figure 10: Jensen-Shannon divergence between predicted and true label proportions, when varying document length, of various losses generated by a Tsallis entropy.

9.2 Dependency parsing experiments

We next compare Fenchel-Young losses for structured output prediction, as formulated in §7, for non-projective *dependency parsing*. The application consists of predicting the directed tree of grammatical dependencies between words in a sentence (Jurafsky and Martin, 2018, Chapter 14). We tackle it by learning structured models over the arborescence polytope, using the structured SVM, CRF, and SparseMAP losses (Niculae et al., 2018), which we saw to all be instances of Fenchel-Young loss. Training data, as in the CoNLL 2017 shared task (Zeman et al., 2017), comes from the Universal Dependency project (Nivre et al., 2016). To isolate the effect of the loss, we use the provided gold tokenization and part-of-speech tags. We follow closely the bidirectional LSTM arc-factored parser of Kiperwasser and Goldberg (2016), using the same model configuration; the only exception is not using externally pretrained embeddings. Parameters are trained using Adam (Kingma and Ba,

Table 5: Unlabeled attachment accuracy scores for dependency parsing, using a bi-LSTM model (Kiperwasser and Goldberg, 2016). For context, we include the scores of the CoNLL 2017 UDPipe baseline (Straka and Straková, 2017).

Loss	en	zh	vi	ro	ja
Structured SVM	87.02	81.94	69.42	87.58	96.24
CRF	86.74	83.18	69.10	87.13	96.09
SparseMAP	86.90	84.03	69.71	87.35	96.04
UDPipe baseline	87.68	82.14	69.63	87.36	95.94

2015), tuning the learning rate on the grid $\{.5, 1, 2, 4, 8\} \times 10^{-3}$, expanded by a factor of 2 if the best model is at either end.

We experiment with 5 languages, diverse both in terms of family and in terms of the amount of training data (ranging from 1,400 sentences for Vietnamese to 12,525 for English). Test set results (Table 5) indicate that SparseMAP outperforms the CRF loss and is competitive with the structured SVM loss, outperforming it substantially on Chinese. The Fenchel-Young perspective sheds further light onto the empirical results, as the SparseMAP loss and the structured SVM (unlike the CRF) both enjoy the margin property with margin 1, while both SparseMAP and the CRF loss (unlike structured SVM) are differentiable.

10. Related work

10.1 Proper losses (proper scoring rules)

Proper losses, a.k.a. proper scoring rules, are a well-studied object in statistics (Grünwald and Dawid (2004); Gneiting and Raftery (2007); references therein) and more recently in machine learning (Reid and Williamson, 2010; Williamson et al., 2016). A loss $\ell(\mathbf{p}; \mathbf{y})$, between a ground truth $\mathbf{y} \in \Delta^d$ and a probability forecast $\mathbf{p} \in \Delta^d$, where d is the number of classes, is said to be proper if it is minimized when estimating the true \mathbf{y} . Formally, a proper loss ℓ satisfies

$$\ell(\mathbf{y}; \mathbf{y}) \leq \ell(\mathbf{p}; \mathbf{y}) \quad \forall \mathbf{p}, \mathbf{y} \in \Delta^d.$$

It is strictly proper if the inequality is strict when $\mathbf{p} \neq \mathbf{y}$, implying that it is uniquely minimized by predicting the correct probability. Strictly proper losses induce **Fisher consistent** estimators of probabilities (Williamson et al., 2016). A key result, which dates back to Savage (1971) (see also Gneiting and Raftery (2007)), is that given a regularization function $\Omega: \Delta^d \rightarrow \mathbb{R}$, the function $\ell_\Omega: \Delta^d \times \Delta^d \rightarrow \mathbb{R}_+$ defined by

$$\ell_\Omega(\mathbf{p}; \mathbf{y}) := \langle \nabla \Omega(\mathbf{p}), \mathbf{y} - \mathbf{p} \rangle - \Omega(\mathbf{p})$$

is proper. It is easy to see that $\ell_\Omega(\mathbf{p}; \mathbf{y}) = B_\Omega(\mathbf{y} \parallel \mathbf{p}) - \Omega(\mathbf{y})$, which recovers the well-known Bregman divergence representation of proper losses. For example, using the Gini index

$H(\mathbf{p}) = 1 - \|\mathbf{p}\|^2$ generates the **Brier score** (Brier, 1950)

$$\ell_{-H}(\mathbf{p}; \mathbf{e}_k) = \sum_{i=1}^d ([\![k = i]\!] - p_i)^2,$$

showing that the sparsemax loss and the Brier score share the same generating function.

More generally, while a proper loss ℓ_Ω is related to a **primal-space** Bregman divergence, a Fenchel-Young loss L_Ω can be seen as a **mixed-space** Bregman divergence (§3). This difference has a number of important consequences. First, ℓ_Ω is **not necessarily convex** in \mathbf{p} (Williamson et al. (2016, Proposition 17) show that it is in fact quasi-convex). In contrast, L_Ω is **always** convex in $\boldsymbol{\theta}$. Second, the first argument is **constrained** to Δ^d for ℓ_Ω , while **unconstrained** for L_Ω .

In practice, proper losses are often composed with an **invertible** link function $\boldsymbol{\psi}^{-1}: \mathbb{R}^d \rightarrow \Delta^d$. This form of a loss, $\ell_\Omega(\boldsymbol{\psi}^{-1}(\boldsymbol{\theta}); \mathbf{y})$, is sometimes called composite (Buja et al., 2005; Reid and Williamson, 2010; Williamson et al., 2016). Although the decoupling between loss and link has merits (Reid and Williamson, 2010), the composition of $\ell_\Omega(\cdot; \mathbf{y})$ and $\boldsymbol{\psi}^{-1}(\boldsymbol{\theta})$ is not necessarily convex in $\boldsymbol{\theta}$. The **canonical link function** (Buja et al., 2005) of ℓ_Ω is a link function that ensures the convexity of $\ell_\Omega(\boldsymbol{\psi}^{-1}(\boldsymbol{\theta}); \mathbf{y})$ in $\boldsymbol{\theta}$. It also plays a key role in generalized linear models (Nelder and Baker, 1972; McCullagh and Nelder, 1989). Following Proposition 3, when Ω is Legendre type, we obtain

$$L_\Omega(\boldsymbol{\theta}; \mathbf{y}) = B_\Omega(\mathbf{y} \|\hat{\mathbf{y}}_\Omega(\boldsymbol{\theta})) = \ell_\Omega(\hat{\mathbf{y}}_\Omega(\boldsymbol{\theta}); \mathbf{y}) + \Omega(\mathbf{y}).$$

Thus, in this case, Fenchel-Young losses and proper composite losses coincide up to the constant term $\Omega(\mathbf{y})$ (which vanishes if $\mathbf{y} = \mathbf{e}_k$ and Ω satisfies assumption A.1), with $\boldsymbol{\psi}^{-1} = \hat{\mathbf{y}}_\Omega$ the canonical inverse link function. Fenchel-Young losses, however, require neither invertible link nor Legendre type assumptions, allowing to express losses (e.g., hinge or sparsemax) that are not expressible in composite form. Moreover, as seen in §5, a Legendre-type Ω precisely precludes sparse probability distributions and losses enjoying a margin.

Other related losses. Nock and Nielsen (2009) proposed a binary classification loss construction. Technically, their loss is based on the Legendre transformation (a subset of Fenchel conjugate functions), precluding non-invertible mappings. Masnadi-Shirazi (2011) studied the Bayes consistency of related binary classification loss functions. Duchi et al. (2018, Proposition 3) derived the multi-class loss (14), a special case of Fenchel-Young loss over the probability simplex, and showed (Proposition 4) that any strictly concave generalized entropy generates a classification-calibrated loss. Amid and Warmuth (2017) proposed a different family of losses based on the Tsallis divergence, to interpolate between convex and non-convex losses, for robustness to label noise. Finally, several works have explored connections between another divergence, the f -divergence (Csiszár, 1975), and surrogate loss functions (Nguyen et al., 2009; Reid and Williamson, 2011; Garcia-Garcia and Williamson, 2012; Duchi et al., 2018).

10.2 Fenchel duality in machine learning

In this paper, we make extensive use of Fenchel duality to derive loss functions. Fenchel duality has also played a key role in several machine learning studies before. It was used to provide a unifying perspective on convex empirical risk minimization and representer theorems (Rifkin and Lippert, 2007). It was also used for deriving regret bounds (Shalev-Shwartz and Singer, 2007; Shalev-Shwartz and Kakade, 2009), risk bounds (Kakade et al., 2009), and unified analyses of boosting (Shen and Li, 2010; Shalev-Shwartz and Singer, 2010; Telgarsky, 2012). However, none of these works propose loss functions, as we do. More closely related to our work are smoothing techniques (Nesterov, 2005; Beck and Teboulle, 2012), which have been used extensively to create smoothed versions of existing losses (Shalev-Shwartz and Zhang, 2016). However, these techniques were applied on a per-loss basis and were not connected to an induced probability distribution. In contrast, we propose a generic loss construction, with clear links between smoothing / regularization and the probability distribution produced by regularized prediction functions.

10.3 Approximate inference with conditional gradient algorithms

In this paper, we suggest conditional gradient (CG) algorithms as a powerful tool for computing regularized prediction functions and Fenchel-Young losses over potentially complex output domains. CG algorithms have been used before to compute an approximate solution to intractable marginal inference problems (Belanger et al., 2013; Krishnan et al., 2015) or to sample from an intractable distribution so as to approximate its mean (Bach et al., 2012; Lacoste-Julien et al., 2015). When combined with reproducing kernel Hilbert spaces (RKHS), these ideas are closely related maximum mean discrepancy (Gretton et al., 2012).

11. Conclusion

We showed that the notion of output regularization and Fenchel duality provide simple core principles, unifying many existing loss functions, and allowing to create useful new ones easily, on a large spectrum of tasks. We characterized a tight connection between sparse distributions and losses with a separation margin, and showed that these losses are precisely the ones that cannot be written in proper composite loss form. We established the computational tools to efficiently learn with Fenchel-Young losses, whether in unstructured or structured settings. We expect that this groundwork will enable the creation of many more novel losses in the future, by exploring other convex polytopes and regularizers.

Acknowledgments

MB thanks Arthur Mensch, Gabriel Peyré and Marco Cuturi for numerous fruitful discussions and Tim Vieira for introducing him to generalized exponential families. This work was supported by the European Research Council (ERC StG DeepSPIN 758969) and by

the Fundação para a Ciência e Tecnologia through contracts UID/EEA/50008/2013 and CMUPERI/TIC/0046/2014 (GoLocal).

A. Additional materials

A.1 Computing proximity operators by root finding

In this section, we first show how the proximity operator of uniformly separable generalized entropies reduces to unidimensional root finding. We then illustrate the case of two important entropies. The proximity operator is defined as

$$\begin{aligned} \text{prox}_{\frac{1}{\sigma}\mathsf{H}}(\mathbf{x}) &= \underset{\mathbf{p} \in \Delta^d}{\operatorname{argmin}} \frac{\sigma}{2} \|\mathbf{p} - \mathbf{x}\|^2 - \mathsf{H}(\mathbf{p}) \\ &= \underset{\mathbf{p} \in \Delta^d}{\operatorname{argmax}} \langle \mathbf{p}, \mathbf{x} \rangle + \frac{1}{\sigma} \mathsf{H}(\mathbf{p}) - \frac{1}{2} \|\mathbf{p}\|^2 \\ &= \underset{\mathbf{p} \in \Delta^d}{\operatorname{argmax}} \langle \mathbf{p}, \mathbf{x} \rangle - \sum_{i=1}^d g(p_i), \end{aligned}$$

where

$$g(t) := \frac{1}{2}t^2 - \frac{1}{\sigma}h(t).$$

Note that $g(t)$ is strictly convex even when h is only concave. We may thus apply Proposition 9 to compute $\text{prox}_{\frac{1}{\sigma}\mathsf{H}}$. The ingredients necessary for using Algorithm 1 are $g'(t) = t - \frac{1}{\sigma}h'(t)$, and its inverse $(g')^{-1}(x)$. We next derive closed-form expressions for the inverse for two important entropies.

Shannon entropy. We have

$$g(t) := \frac{t^2}{2} + \frac{t \log t}{\sigma} \quad \text{and} \quad g'(t) = t + \frac{1 + \log t}{\sigma}.$$

To compute the inverse, $(g')^{-1}(x) = t$, we solve for t satisfying

$$\frac{1 + \log t}{\sigma} + t = x.$$

Making the change of variable $s := \sigma t$ gives

$$s + \log s = \sigma x - 1 + \log \sigma \iff s = \omega(\sigma x - 1 + \log \sigma).$$

Therefore, $(g')^{-1}(x) = \frac{1}{\sigma}\omega(\sigma x - 1 + \log \sigma)$, where ω denotes the Wright omega function.

Tsallis entropy with $\alpha = 1.5$. Up to a constant term, we have

$$g(t) = \frac{t^2}{2} + \frac{4}{3\sigma}t^{3/2} \quad \text{and} \quad g'(t) = t + \frac{2}{\sigma}\sqrt{t}.$$

To obtain the inverse, $(g')^{-1}(x) = t$ we seek t satisfying

$$\frac{2}{\sigma}\sqrt{t} + t = x \iff \sqrt{t} = \frac{\sigma}{2}(x - t).$$

If $t > x$ there are no solutions, otherwise we may square both sides, yielding

$$\frac{\sigma^2 t^2}{4} - t \left(\frac{\sigma^2 x}{2} + 1 \right) + \frac{\sigma^2 x^2}{4} = 0.$$

The discriminant is

$$\Delta = 1 + \sigma^2 x > 0,$$

resulting in two solutions

$$t_{\pm} = x + \frac{2}{\sigma^2} \left(1 \pm \sqrt{1 + \sigma^2 x} \right).$$

However, $t_+ > x$, therefore

$$(g')^{-1}(x) = x + \frac{2}{\sigma^2} \left(1 - \sqrt{1 + \sigma^2 x} \right).$$

A.2 Loss ‘‘Fenchel-Youngization’’


Not all loss functions proposed in the literature can be written in Fenchel-Young loss form. In this section, we present a natural method to approximate (in a sense we will clarify) any loss $\ell: \mathbb{R}^d \times \{\mathbf{e}_i\}_{i=1}^d \rightarrow \mathbb{R}_+$ with a Fenchel-Young loss. This has two main advantages. First, the resulting loss is convex even when ℓ is not. Second, the associated regularized prediction function can be used for probabilistic prediction even if ℓ is not probabilistic.

From loss to entropy. Equations (16) and (17), which relate entropy and conditional Bayes risk, suggest a reverse construction from entropy to loss function. This direction has been explored in previous works (Grünwald and Dawid, 2004; Duchi et al., 2018). The entropy H_{ℓ} generated from ℓ is defined as follows:

$$H_{\ell}(\mathbf{p}) := \inf_{\boldsymbol{\theta} \in \mathbb{R}^d} \mathbb{E}_{\mathbf{p}}[\ell(\boldsymbol{\theta}; Y)] = \inf_{\boldsymbol{\theta} \in \mathbb{R}^d} \sum_{i=1}^d p_i \ell(\boldsymbol{\theta}; \mathbf{e}_i). \quad (34)$$

This is the infimum of a linear and thus concave function of \mathbf{p} . Hence, by Danskin’s theorem (Danskin, 1966; Bertsekas, 1999, Proposition B.25), $H_{\ell}(\mathbf{p})$ is concave (note that this is true even if ℓ is not convex). As an example, Duchi et al. (2018, Example 1) show that the entropy associated with the zero-one loss is $H_{\ell}(\mathbf{p}) = 1 - \max_j p_j$. As we discussed in §4.3, this entropy is known as the Berger-Parker dominance index (Berger and Parker, 1970), and is a special case of norm entropy (20).

The following new proposition shows how to compute the entropy of pairwise hinge losses.


 Figure 11: Contour comparison of Gini index H_2^T and pairwise hinge entropies H_ℓ .

Proposition 10 *Entropy generated by pairwise hinge losses*

Let $\ell(\boldsymbol{\theta}; \mathbf{e}_k) := \sum_{j \neq k} \phi(\theta_j - \theta_k)$ and $\tau(\mathbf{p}) := \min_j 1 - p_j$. Then, for all $\mathbf{p} \in \Delta^d$:

1. If $\phi(t)$ is the hinge function $[1+t]_+$ then $H_\ell(\mathbf{p}) = \tau(\mathbf{p})d$.
2. If $\phi(t)$ is the smoothed hinge function (48), then $H_\ell(\mathbf{p}) = \frac{-\tau(\mathbf{p})^2}{2} \sum_{j=1}^d \frac{1}{1-p_j} + \tau(\mathbf{p})d$.
3. If $\phi(t)$ is the squared hinge function $\frac{1}{2}[1+t]_+^2$, then $H_\ell(\mathbf{p}) = \frac{\frac{1}{2}d^2}{\sum_{j=1}^d 1/(1-p_j)}$.

See Appendix B.9 for a proof. The first one recovers (Duchi et al., 2018, Example 5) with a simpler proof. The last two are new. These entropies are illustrated in Figure 11. For the last two choices, H_ℓ is strictly concave over Δ^d and the associated prediction function $\hat{\mathbf{y}}_{-H_\ell}(\boldsymbol{\theta})$ is typically sparse, although we are not aware of a closed-form expression.

And back to Fenchel-Young loss. A natural idea is then to construct a Fenchel-Young from H_ℓ . The resulting loss is convex by construction. Let us assume for simplicity that ℓ achieves its minimum at 0, implying $H_\ell(\mathbf{e}_k) = 0$. Combining (14) and (34), we obtain

$$\begin{aligned} L_{-H_\ell}(\boldsymbol{\theta}; \mathbf{e}_k) &= (-H_\ell)^*(\boldsymbol{\theta}) - \theta_k \\ &= \max_{\mathbf{p} \in \Delta^d} \langle \boldsymbol{\theta}, \mathbf{p} \rangle + H_\ell(\mathbf{p}) - \theta_k \\ &= \max_{\mathbf{p} \in \Delta^d} \min_{\boldsymbol{\theta}' \in \mathbb{R}^d} \sum_{i=1}^d p_i (\ell(\boldsymbol{\theta}'; \mathbf{e}_i) + \theta_i) - \theta_k. \end{aligned}$$

By weak duality, we have for every (possibly non-convex) ℓ

$$\begin{aligned} L_{-H_\ell}(\boldsymbol{\theta}; \mathbf{e}_k) &\leq \min_{\boldsymbol{\theta}' \in \mathbb{R}^d} \max_{\mathbf{p} \in \Delta^d} \sum_{i=1}^d p_i (\ell(\boldsymbol{\theta}'; \mathbf{e}_i) + \theta_i) - \theta_k \\ &= \min_{\boldsymbol{\theta}' \in \mathbb{R}^d} \max_{i \in [d]} \ell(\boldsymbol{\theta}'; \mathbf{e}_i) + \theta_i - \theta_k. \end{aligned}$$

We can see $(\boldsymbol{\theta}; \mathbf{e}_k) \mapsto \max_{i \in [d]} \ell(\boldsymbol{\theta}'; \mathbf{e}_i) + \theta_i - \theta_k$ as a family of cost-augmented hinge losses (Tsochantaridis et al., 2005) parametrized by $\boldsymbol{\theta}'$. Hence, $L_{-H_\ell}(\boldsymbol{\theta}; \mathbf{e}_k)$ is upper-bounded by the tightest loss in this family.

If ℓ is convex, strong duality holds and we obtain equality

$$L_{-\text{H}\ell}(\boldsymbol{\theta}; \mathbf{e}_k) = \min_{\boldsymbol{\theta}' \in \mathbb{R}^d} \max_{i \in [d]} \ell(\boldsymbol{\theta}'; \mathbf{e}_i) + \theta_i - \theta_k.$$

Margin. We can easily upper-bound the margin of $L_{-\text{H}\ell}$. Indeed, assuming $\ell(\boldsymbol{\theta}; \mathbf{e}_i)$ is upper-bounded by m for all $\boldsymbol{\theta}$ and $i \in [d]$, we get

$$0 \leq L_{-\text{H}\ell}(\boldsymbol{\theta}; \mathbf{e}_k) \leq \max_{i \in [d]} \ell(\boldsymbol{\theta}; \mathbf{e}_i) + \theta_i - \theta_k \leq \max_{i \in [d]} m + \theta_i - \theta_k.$$

Hence, from Definition 4, we have $\text{margin}(L_{-\text{H}\ell}) \leq m$. Note that if $\text{margin}(L_{-\text{H}\ell}) \leq \text{margin}(\ell)$, then $L_{-\text{H}\ell}(\boldsymbol{\theta}; \mathbf{e}_k) = 0 \Rightarrow \ell(\boldsymbol{\theta}; \mathbf{e}_k) = 0$.

B. Proofs

B.1 Proof of Proposition 1

Effect of a permutation. Let $\Omega(\boldsymbol{\mu})$ be symmetric. We first prove that Ω^* is symmetric as well. Indeed, we have

$$\Omega^*(\mathbf{P}\boldsymbol{\theta}) = \sup_{\boldsymbol{\mu} \in \text{dom}(\Omega)} (\mathbf{P}\boldsymbol{\theta})^\top \boldsymbol{\mu} - \Omega(\boldsymbol{\mu}) = \sup_{\boldsymbol{\mu} \in \text{dom}(\Omega)} \boldsymbol{\theta}^\top \mathbf{P}^\top \boldsymbol{\mu} - \Omega(\mathbf{P}^\top \boldsymbol{\mu}) = \Omega^*(\boldsymbol{\theta}).$$

The last equality was obtained by a change of variable $\boldsymbol{\mu}' = \mathbf{P}^\top \boldsymbol{\mu}$, from which $\boldsymbol{\mu}$ is recovered as $\boldsymbol{\mu} = \mathbf{P}\boldsymbol{\mu}'$, which proves $\nabla \Omega^*(\mathbf{P}\boldsymbol{\mu}) = \mathbf{P} \nabla \Omega^*(\boldsymbol{\mu})$.

Order preservation. Since Ω^* is convex, the gradient operator $\nabla \Omega^*$ is monotone, i.e.,

$$(\boldsymbol{\theta}' - \boldsymbol{\theta})^\top (\boldsymbol{\mu}' - \boldsymbol{\mu}) \geq 0$$

for any $\boldsymbol{\theta}, \boldsymbol{\theta}' \in \mathbb{R}^d$, $\boldsymbol{\mu} = \nabla \Omega^*(\boldsymbol{\theta})$ and $\boldsymbol{\mu}' = \nabla \Omega^*(\boldsymbol{\theta}')$. Let $\boldsymbol{\theta}'$ be obtained from $\boldsymbol{\theta}$ by swapping two coordinates, i.e., $\theta'_j = \theta_i$, $\theta'_i = \theta_j$, and $\theta'_k = \theta_k$ for any $k \notin \{i, j\}$. Then, since Ω is symmetric, we obtain:

$$2(\theta_j - \theta_i)(\mu_j - \mu_i) \geq 0,$$

which implies $\theta_i > \theta_j \Rightarrow \mu_i \geq \mu_j$ and $\mu_i > \mu_j \Rightarrow \theta_i \geq \theta_j$. To fully prove the claim, we need to show that the last inequality is strict: to do this, we simply invoke $\nabla \Omega^*(\mathbf{P}\boldsymbol{\mu}) = \mathbf{P} \nabla \Omega^*(\boldsymbol{\mu})$ with a matrix \mathbf{P} that permutes i and j , from which we must have $\theta_i = \theta_j \Rightarrow \mu_i = \mu_j$.

Approximation error. Let $\mathbf{y}^* := \text{argmax}_{\mathbf{y} \in \mathcal{Y}} \langle \boldsymbol{\theta}, \mathbf{y} \rangle$ and $\boldsymbol{\mu}^* := \text{argmax}_{\boldsymbol{\mu} \in \text{dom}(\Omega)} \langle \boldsymbol{\theta}, \boldsymbol{\mu} \rangle - \Omega(\boldsymbol{\mu})$. If f is a differentiable γ -strongly convex function with unique minimizer $\boldsymbol{\mu}^*$, then

$$\frac{\gamma}{2} \|\mathbf{y} - \boldsymbol{\mu}^*\|_2^2 \leq f(\mathbf{y}) - f(\boldsymbol{\mu}^*) \quad \forall \mathbf{y} \in \text{dom}(f).$$

We assume that $\mathcal{Y} \subseteq \text{dom}(\Omega)$ so it suffices to upper-bound $f(\mathbf{y}^*) - f(\boldsymbol{\mu}^*)$, where $f(\boldsymbol{\mu}) := \Omega(\boldsymbol{\mu}) - \langle \boldsymbol{\theta}, \boldsymbol{\mu} \rangle$. Since $L \leq \Omega(\boldsymbol{\mu}) \leq U$ for all $\boldsymbol{\mu} \in \text{dom}(\Omega)$, we have $\langle \boldsymbol{\theta}, \boldsymbol{\mu}^* \rangle \geq -f(\boldsymbol{\mu}^*) + L$ and $-f(\mathbf{y}^*) + U \geq \langle \boldsymbol{\theta}, \mathbf{y}^* \rangle$. Together with $\langle \boldsymbol{\theta}, \mathbf{y}^* \rangle \geq \langle \boldsymbol{\theta}, \boldsymbol{\mu}^* \rangle$, this implies $f(\mathbf{y}^*) - f(\boldsymbol{\mu}^*) \leq U - L$. Therefore, $\frac{1}{2} \|\mathbf{y}^* - \boldsymbol{\mu}^*\|^2 \leq \frac{U-L}{\gamma}$.

Temperature scaling & constant invariance. These immediately follow from properties of the argmax operator.

B.2 Proof of Proposition 3

We set $\Omega := \Psi + I_{\mathcal{C}}$.

Bregman projections. If Ψ is Legendre type, then $\nabla\Psi(\nabla\Psi^*(\boldsymbol{\theta})) = \boldsymbol{\theta}$ for all $\boldsymbol{\theta} \in \text{int}(\text{dom}(\Psi^*))$. Using this and our assumption that $\text{dom}(\Psi^*) = \mathbb{R}^d$, we get for all $\boldsymbol{\theta} \in \mathbb{R}^d$:

$$B_{\Psi}(\mathbf{p} \parallel \nabla\Psi^*(\boldsymbol{\theta})) = \Psi(\mathbf{p}) - \langle \boldsymbol{\theta}, \mathbf{p} \rangle + \langle \boldsymbol{\theta}, \nabla\Psi^*(\boldsymbol{\theta}) \rangle - \Psi(\nabla\Psi^*(\boldsymbol{\theta})). \quad (36)$$

The last two terms are independent of \mathbf{p} and therefore

$$\hat{\mathbf{y}}_{\Omega}(\boldsymbol{\theta}) = \underset{\mathbf{p} \in \mathcal{C}}{\text{argmax}} \langle \boldsymbol{\theta}, \mathbf{p} \rangle - \Psi(\mathbf{p}) = \underset{\mathbf{p} \in \mathcal{C}}{\text{argmin}} B_{\Psi}(\mathbf{p} \parallel \nabla\Psi^*(\boldsymbol{\theta})),$$

where $\mathcal{C} \subseteq \text{dom}(\Psi)$. The r.h.s. is the Bregman projection of $\nabla\Psi^*(\boldsymbol{\theta})$ onto \mathcal{C} .

Difference of Bregman divergences. Let $\mathbf{p} = \hat{\mathbf{y}}_{\Omega}(\boldsymbol{\theta})$. Using (36), we obtain

$$\begin{aligned} B_{\Psi}(\mathbf{y} \parallel \nabla\Psi^*(\boldsymbol{\theta})) - B_{\Psi}(\mathbf{p} \parallel \nabla\Psi^*(\boldsymbol{\theta})) &= \Psi(\mathbf{y}) - \langle \boldsymbol{\theta}, \mathbf{y} \rangle + \langle \boldsymbol{\theta}, \mathbf{p} \rangle - \Psi(\mathbf{p}) \\ &= \Omega(\mathbf{y}) - \langle \boldsymbol{\theta}, \mathbf{y} \rangle + \Omega^*(\boldsymbol{\theta}) \\ &= L_{\Omega}(\boldsymbol{\theta}; \mathbf{y}), \end{aligned} \quad (37)$$

where we assumed $\mathbf{y} \in \mathcal{C}$ and $\mathcal{C} \subseteq \text{dom}(\Psi)$, implying $\Psi(\mathbf{y}) = \Omega(\mathbf{y})$.

If $\mathcal{C} = \text{dom}(\Psi)$ (i.e., $\Omega = \Psi$), then $\mathbf{p} = \nabla\Psi^*(\boldsymbol{\theta})$ and $B_{\Psi}(\mathbf{p} \parallel \nabla\Psi^*(\boldsymbol{\theta})) = 0$. We thus get the **composite form** of Fenchel-Young losses

$$B_{\Omega}(\mathbf{y} \parallel \nabla\Omega^*(\boldsymbol{\theta})) = B_{\Omega}(\mathbf{y} \parallel \hat{\mathbf{y}}_{\Omega}(\boldsymbol{\theta})) = L_{\Omega}(\boldsymbol{\theta}; \mathbf{y}).$$

Bound. Let $\mathbf{p} = \hat{\mathbf{y}}_{\Omega}(\boldsymbol{\theta})$. Since \mathbf{p} is the Bregman projection of $\nabla\Psi^*(\boldsymbol{\theta})$ onto \mathcal{C} , we can use the well-known Pythagorean theorem for Bregman divergences (see, e.g., Banerjee et al. (2005, Appendix A)) to obtain for all $\mathbf{y} \in \mathcal{C} \subseteq \text{dom}(\Psi)$:

$$B_{\Psi}(\mathbf{y} \parallel \mathbf{p}) + B_{\Psi}(\mathbf{p} \parallel \nabla\Psi^*(\boldsymbol{\theta})) \leq B_{\Psi}(\mathbf{y} \parallel \nabla\Psi^*(\boldsymbol{\theta})).$$

Using (37), we obtain for all $\mathbf{y} \in \mathcal{C}$:

$$0 \leq B_{\Psi}(\mathbf{y} \parallel \mathbf{p}) \leq L_{\Omega}(\boldsymbol{\theta}; \mathbf{y}).$$

Since Ω is a l.s.c. proper convex function, from Proposition 2, we immediately get

$$\mathbf{p} = \mathbf{y} \Leftrightarrow L_{\Omega}(\boldsymbol{\theta}; \mathbf{y}) = 0 \Leftrightarrow B_{\Psi}(\mathbf{y} \parallel \mathbf{p}) = 0.$$

B.3 Proof of Proposition 4

The two facts stated in Proposition 4 (H is always non-negative and maximized by the uniform distribution) follow directly from Jensen's inequality. Indeed, for all $\mathbf{p} \in \Delta^d$:

- $H(\mathbf{p}) \geq \sum_{j=1}^d p_j H(\mathbf{e}_j) = 0$;
- $H(\mathbf{1}/d) = H\left(\sum_{\mathbf{P} \in \mathcal{P}} \frac{1}{d!} \mathbf{P}\mathbf{p}\right) \geq \sum_{\mathbf{P} \in \mathcal{P}} \frac{1}{d!} H(\mathbf{P}\mathbf{p}) = H(\mathbf{p})$,

where \mathcal{P} is the set of $d \times d$ permutation matrices. Strict concavity ensures that $\mathbf{p} = \mathbf{1}/d$ is the unique maximizer.

B.4 Proof of Proposition 5

We start by proving the following lemma.

Lemma 1 *Let H satisfy assumptions A.1–A.3. Then:*

1. *We have $\boldsymbol{\theta} \in \partial(-H)(\mathbf{e}_k)$ iff $\theta_k = (-H)^*(\boldsymbol{\theta})$. That is:*

$$\partial(-H)(\mathbf{e}_k) = \{\boldsymbol{\theta} \in \mathbb{R}^d : \theta_k \geq \langle \boldsymbol{\theta}, \mathbf{p} \rangle + H(\mathbf{p}), \forall \mathbf{p} \in \Delta^d\}.$$

2. *If $\boldsymbol{\theta} \in \partial(-H)(\mathbf{e}_k)$, then, we also have $\boldsymbol{\theta}' \in \partial(-H)(\mathbf{e}_k)$ for any $\boldsymbol{\theta}'$ such that $\theta'_k = \theta_k$ and $\theta'_i \leq \theta_i$, for all $i \neq k$.*

Proof of the lemma: Let $\Omega = -H$. From Proposition 1 (order preservation), we can consider $\partial\Omega(\mathbf{e}_1)$ without loss of generality, in which case any $\boldsymbol{\theta} \in \partial\Omega(\mathbf{e}_1)$ satisfies $\theta_1 = \max_j \theta_j$. We have $\boldsymbol{\theta} \in \partial\Omega(\mathbf{e}_1)$ iff $\Omega(\mathbf{e}_1) = \langle \boldsymbol{\theta}, \mathbf{e}_1 \rangle - \Omega^*(\boldsymbol{\theta}) = \theta_1 - \Omega^*(\boldsymbol{\theta})$. Since $\Omega(\mathbf{e}_1) = 0$, we must have $\theta_1 = \Omega^*(\boldsymbol{\theta}) \geq \sup_{\mathbf{p} \in \Delta^d} \langle \boldsymbol{\theta}, \mathbf{p} \rangle - \Omega(\mathbf{p})$, which proves part 1. To see 2, note that we have $\theta'_k = \theta_k \geq \langle \boldsymbol{\theta}, \mathbf{p} \rangle - \Omega(\mathbf{p}) \geq \langle \boldsymbol{\theta}', \mathbf{p} \rangle - \Omega(\mathbf{p})$, for all $\mathbf{p} \in \Delta^d$, from which the result follows. \blacksquare

We now proceed to the proof of Proposition 5. Let $\Omega = -H$, and suppose that L_Ω has the separation margin property. Then, $\boldsymbol{\theta} = m\mathbf{e}_1$ satisfies the margin condition $\theta_1 \geq m + \max_{j \neq 1} \theta_j$, hence $L_\Omega(m\mathbf{e}_1, \mathbf{e}_1) = 0$. From the first part of Proposition 2, this implies $m\mathbf{e}_1 \in \partial\Omega(\mathbf{e}_1)$.

Conversely, let us assume that $m\mathbf{e}_1 \in \partial\Omega(\mathbf{e}_1)$. From the second part of Lemma 1, this implies that $\boldsymbol{\theta} \in \partial\Omega(\mathbf{e}_1)$ for any $\boldsymbol{\theta}$ such that $\theta_1 = m$ and $\theta_i \leq 0$ for all $i \geq 2$; and more generally we have $\boldsymbol{\theta} + c\mathbf{1} \in \partial\Omega(\mathbf{e}_1)$. That is, any $\boldsymbol{\theta}$ with $\theta_1 \geq m + \max_{i \neq 1} \theta_i$ satisfies $\boldsymbol{\theta} \in \partial\Omega(\mathbf{e}_1)$. From Proposition 2, this is equivalent to $L_\Omega(\boldsymbol{\theta}; \mathbf{e}_1) = 0$.

Let us now determine the margin of L_Ω , i.e., the smallest m such that $m\mathbf{e}_1 \in \partial\Omega(\mathbf{e}_1)$. From Lemma 1, this is equivalent to $m \geq mp_1 - \Omega(\mathbf{p})$ for any $\mathbf{p} \in \Delta^d$, i.e., $\frac{-\Omega(\mathbf{p})}{1-p_1} \leq m$. Note that by Proposition 1 the “most competitive” \mathbf{p} 's are sorted as \mathbf{e}_1 , so we may write $p_1 = \|\mathbf{p}\|_\infty$ without loss of generality. The margin of L_Ω is the smallest possible such margin, given by (22).

B.5 Proof of Proposition 6

Let us start by showing that conditions 1 and 2 are equivalent. To show that $2 \Rightarrow 1$, take an arbitrary $\mathbf{p} \in \Delta^d$. From Fenchel-Young duality and the Danskin's theorem, we have that $\nabla(-\mathbf{H})^*(\boldsymbol{\theta}) = \mathbf{p} \Rightarrow \boldsymbol{\theta} \in \partial(-\mathbf{H})(\mathbf{p})$, which implies the subdifferential set is non-empty everywhere in the simplex. Let us now prove that $1 \Rightarrow 2$. Let $\Omega = -\mathbf{H}$, and assume that Ω has non-empty subdifferential everywhere in Δ^d . We need to show that for any $\mathbf{p} \in \Delta^d$, there is some $\boldsymbol{\theta} \in \mathbb{R}^d$ such that $\mathbf{p} \in \operatorname{argmin}_{\mathbf{p}' \in \Delta^d} \Omega(\mathbf{p}') - \langle \boldsymbol{\theta}, \mathbf{p}' \rangle$. The Lagrangian associated with this minimization problem is:

$$\mathcal{L}(\mathbf{p}, \boldsymbol{\mu}, \lambda) = \Omega(\mathbf{p}) - \langle \boldsymbol{\theta} + \boldsymbol{\mu}, \mathbf{p} \rangle + \lambda(\mathbf{1}^\top \mathbf{p} - 1).$$

The KKT conditions are:

$$\begin{cases} 0 \in \partial_{\mathbf{p}} \mathcal{L}(\mathbf{p}, \boldsymbol{\mu}, \lambda) = \partial \Omega(\mathbf{p}) - \boldsymbol{\theta} - \boldsymbol{\mu} + \lambda \mathbf{1} \\ \langle \mathbf{p}, \boldsymbol{\mu} \rangle = 0 \\ \mathbf{p} \in \Delta^d, \boldsymbol{\mu} \geq 0. \end{cases}$$

For a given $\mathbf{p} \in \Delta^d$, we seek $\boldsymbol{\theta}$ such that $(\mathbf{p}, \boldsymbol{\mu}, \lambda)$ are a solution to the KKT conditions for some $\boldsymbol{\mu} \geq 0$ and $\lambda \in \mathbb{R}$.

We will show that such $\boldsymbol{\theta}$ exists by simply choosing $\boldsymbol{\mu} = \mathbf{0}$ and $\lambda = 0$. Those choices are dual feasible and guarantee that the slackness complementary condition is satisfied. In this case, we have from the first condition that $\boldsymbol{\theta} \in \partial \Omega(\mathbf{p})$. From the assumption that Ω has non-empty subdifferential in all the simplex, we have that for any $\mathbf{p} \in \Delta^d$ we can find a $\boldsymbol{\theta} \in \mathbb{R}^d$ such that $(\mathbf{p}, \boldsymbol{\theta})$ are a dual pair, i.e., $\mathbf{p} = \nabla \Omega^*(\boldsymbol{\theta})$, which proves that $\nabla \Omega^*(\mathbb{R}^d) = \Delta^d$.

Next, we show that condition $1 \Rightarrow 3$. Since $\partial(-\mathbf{H})(\mathbf{p}) \neq \emptyset$ everywhere in the simplex, we can take an arbitrary $\boldsymbol{\theta} \in \partial(-\mathbf{H})(\mathbf{e}_k)$. From Lemma 1, item 2, we have that $\boldsymbol{\theta}' \in \partial(-\mathbf{H})(\mathbf{e}_k)$ for $\theta'_k = \theta_k$ and $\theta'_j = \min_{\ell} \theta_\ell$; since $(-\mathbf{H})^*$ is shift invariant, we can without loss of generality have $\boldsymbol{\theta}' = m\mathbf{e}_k$ for some $m > 0$, which implies from Proposition 5 that L_Ω has a margin.

Let us show that, if $-\mathbf{H}$ is separable, then $3 \Rightarrow 1$, which establishes equivalence between all conditions 1, 2, and 3. From Proposition 5, the existing of a separation margin implies that there is some m such that $m\mathbf{e}_k \in \partial(-\mathbf{H})(\mathbf{e}_k)$. Let $\mathbf{H}(\mathbf{p}) = \sum_{i=1}^d h(p_i)$, with $h : [0, 1] \rightarrow \mathbb{R}_+$ concave. Due to assumption A.1, h must satisfy $h(0) = h(1) = 0$. Without loss of generality, suppose $\mathbf{p} = [\tilde{\mathbf{p}}; \mathbf{0}_k]$, where $\tilde{\mathbf{p}} \in \operatorname{relint}(\Delta^{d-k})$ and $\mathbf{0}_k$ is a vector with k zeros. We will see that there is a vector $\mathbf{g} \in \mathbb{R}^d$ such that $\mathbf{g} \in \partial(-\mathbf{H})(\mathbf{p})$, i.e., satisfying

$$-\mathbf{H}(\mathbf{p}') \geq -\mathbf{H}(\mathbf{p}) + \langle \mathbf{g}, \mathbf{p}' - \mathbf{p} \rangle, \quad \forall \mathbf{p}' \in \Delta^d. \quad (39)$$

Since $\tilde{\mathbf{p}} \in \operatorname{relint}(\Delta^{d-k})$, we have $\tilde{p}_i \in]0, 1[$ for $i \in \{1, \dots, d-k\}$, hence $\partial(-h)(\tilde{p}_i)$ must be nonempty, since $-h$ is convex and $]0, 1[$ is an open set. We show that the following $\mathbf{g} = (g_1, \dots, g_d) \in \mathbb{R}^d$ is a subgradient of $-\mathbf{H}$ at \mathbf{p} :

$$g_i = \begin{cases} \partial(-h)(\tilde{p}_i), & i = 1, \dots, d-k \\ m, & i = d-k+1, \dots, d. \end{cases}$$

By definition of subgradient, we have

$$-\psi(p'_i) \geq -\psi(\tilde{p}_i) + \partial(-h)(\tilde{p}_i)(p'_i - \tilde{p}_i), \quad \text{for } i = 1, \dots, d-k. \quad (40)$$

Furthermore, since m upper bounds the separation margin of \mathbf{H} , we have from Proposition 5 that $m \geq \frac{\mathbf{H}([1-p'_i, p'_i, 0, \dots, 0])}{1 - \max\{1-p'_i, p'_i\}} = \frac{h(1-p'_i) + h(p'_i)}{\min\{p'_i, 1-p'_i\}} \geq \frac{h(p'_i)}{p'_i}$ for any $p'_i \in]0, 1]$. Hence, we have

$$-\psi(p'_i) \geq -\psi(0) - m(p'_i - 0), \quad \text{for } i = d - k + 1, \dots, d. \quad (41)$$

Summing all inequalities in Eqs. (40)–(41), we obtain the expression in Eq. (39), which finishes the proof.

B.6 Proof of Proposition 7

Define $\Omega = -\mathbf{H}$. Let us start by writing the margin expression (22) as a unidimensional optimization problem. This is done by noticing that the max-generalized entropy problem constrained to $\max(\mathbf{p}) = 1 - t$ gives $\mathbf{p} = \left[1 - t, \frac{t}{d-1}, \dots, \frac{t}{d-1}\right]$, for $t \in [0, 1 - \frac{1}{d}]$ by a similar argument as the one used in Proposition 4. We obtain:

$$\text{margin}(L_\Omega) = \sup_{t \in [0, 1 - \frac{1}{d}]} \frac{-\Omega\left(\left[1 - t, \frac{t}{d-1}, \dots, \frac{t}{d-1}\right]\right)}{t}.$$

We write the argument above as $A(t) = \frac{-\Omega(\mathbf{e}_1 + t\mathbf{v})}{t}$, where $\mathbf{v} := [-1, \frac{1}{d-1}, \dots, \frac{1}{d-1}]$. We will first prove that A is decreasing in $[0, 1 - \frac{1}{d}]$, which implies that the supremum (and the margin) equals $A(0)$. Note that we have the following expression for the derivative of any function $f(\mathbf{e}_1 + t\mathbf{v})$:

$$(f(\mathbf{e}_1 + t\mathbf{v}))' = \mathbf{v}^\top \nabla f(\mathbf{e}_1 + t\mathbf{v}).$$

Using this fact, we can write the derivative $A'(t)$ as:

$$A'(t) = \frac{-t\mathbf{v}^\top \nabla \Omega(\mathbf{e}_1 + t\mathbf{v}) + \Omega(\mathbf{e}_1 + t\mathbf{v})}{t^2} := \frac{B(t)}{t^2}.$$

In turn, the derivative $B'(t)$ is:

$$\begin{aligned} B'(t) &= -\mathbf{v}^\top \nabla \Omega(\mathbf{e}_1 + t\mathbf{v}) - t(\mathbf{v}^\top \nabla \Omega(\mathbf{e}_1 + t\mathbf{v}))' + \mathbf{v}^\top \nabla \Omega(\mathbf{e}_1 + t\mathbf{v}) \\ &= -t(\mathbf{v}^\top \nabla \Omega(\mathbf{e}_1 + t\mathbf{v}))' \\ &= -t\mathbf{v}^\top \nabla \nabla \Omega(\mathbf{e}_1 + t\mathbf{v})\mathbf{v} \\ &\leq 0, \end{aligned}$$

where we denote by $\nabla \nabla \Omega$ the Hessian of Ω , and used the fact that it is positive semi-definite, due to the convexity of Ω . This implies that B is decreasing, hence for any $t \in [0, 1]$, $B(t) \leq B(0) = \Omega(\mathbf{e}_1) = 0$, where we used the fact $\|\nabla \Omega(\mathbf{e}_1)\| < \infty$, assumed as a condition of Proposition 6. Therefore, we must also have $A'(t) = \frac{B(t)}{t^2} \leq 0$ for any $t \in [0, 1]$, hence A is decreasing, and $\sup_{t \in [0, 1 - \frac{1}{d}]} A(t) = \lim_{t \rightarrow 0^+} A(t)$. By L'Hôpital's rule:

$$\begin{aligned} \lim_{t \rightarrow 0^+} A(t) &= \lim_{t \rightarrow 0^+} (-\Omega(\mathbf{e}_1 + t\mathbf{v}))' \\ &= -\mathbf{v}^\top \nabla \Omega(\mathbf{e}_1) \\ &= \nabla_1 \Omega(\mathbf{e}_1) - \frac{1}{d-1} \sum_{j \geq 2} \nabla_j \Omega(\mathbf{e}_1) \\ &= \nabla_1 \Omega(\mathbf{e}_1) - \nabla_2 \Omega(\mathbf{e}_1), \end{aligned}$$

which proves the first part.

If Ω is separable, then $\nabla_j \Omega(\mathbf{p}) = -h'(p_j)$, in particular $\nabla_1 \Omega(\mathbf{e}_1) = -h'(1)$ and $\nabla_2 \Omega(\mathbf{e}_1) = -h'(0)$, yielding $\text{margin}(L_\Omega) = h'(0) - h'(1)$. Since h is twice differentiable, this equals $-\int_0^1 h''(t)dt$, completing the proof.

B.7 Proof of Proposition 8

We start by proving the following lemma, which generalizes Lemma 1.

Lemma 2 *Let Ω be convex. Then:*

1. We have $\partial\Omega(\mathbf{y}) = \{\boldsymbol{\theta} \in \mathbb{R}^d: \langle \boldsymbol{\theta}, \mathbf{y} \rangle - \Omega(\mathbf{y}) \geq \langle \boldsymbol{\theta}, \boldsymbol{\mu} \rangle - \Omega(\boldsymbol{\mu}), \forall \boldsymbol{\mu} \in \text{conv}(\mathcal{Y})\}$.
2. If $\boldsymbol{\theta} \in \partial\Omega(\mathbf{y})$, then, we also have $\boldsymbol{\theta}' \in \partial\Omega(\mathbf{y})$ for any $\boldsymbol{\theta}'$ such that $\langle \boldsymbol{\theta}' - \boldsymbol{\theta}, \mathbf{y}' \rangle \leq \langle \boldsymbol{\theta}' - \boldsymbol{\theta}, \mathbf{y} \rangle$, for all $\mathbf{y}' \in \mathcal{Y}$.

Proof of the lemma: The first part comes directly from the definition of Fenchel conjugate. To see the second part, note that, if $\boldsymbol{\theta} \in \partial\Omega(\mathbf{y})$, then $\langle \boldsymbol{\theta}', \mathbf{y} \rangle - \Omega(\mathbf{y}) = \langle \boldsymbol{\theta}' - \boldsymbol{\theta}, \mathbf{y} \rangle + \langle \boldsymbol{\theta}, \mathbf{y} \rangle - \Omega(\mathbf{y}) \geq \langle \boldsymbol{\theta}' - \boldsymbol{\theta}, \mathbf{y} \rangle + \langle \boldsymbol{\theta}, \boldsymbol{\mu} \rangle - \Omega(\boldsymbol{\mu})$ for every $\boldsymbol{\mu} \in \text{conv}(\mathcal{Y})$. Let $\boldsymbol{\mu} = \sum_{\mathbf{y}' \in \mathcal{Y}} p(\mathbf{y}') \mathbf{y}'$ for some distribution $\mathbf{p} \in \Delta^{|\mathcal{Y}|}$. Then, we have $\langle \boldsymbol{\theta}' - \boldsymbol{\theta}, \mathbf{y} \rangle + \langle \boldsymbol{\theta}, \boldsymbol{\mu} \rangle - \Omega(\boldsymbol{\mu}) = \langle \boldsymbol{\theta}' - \boldsymbol{\theta}, \mathbf{y} \rangle + \sum_{\mathbf{y}' \in \mathcal{Y}} p(\mathbf{y}') \langle \boldsymbol{\theta}, \mathbf{y}' \rangle - \Omega(\boldsymbol{\mu}) \geq \langle \boldsymbol{\theta}' - \boldsymbol{\theta}, \mathbf{y} \rangle + \sum_{\mathbf{y}' \in \mathcal{Y}} p(\mathbf{y}') (\langle \boldsymbol{\theta}', \mathbf{y}' \rangle - \langle \boldsymbol{\theta}' - \boldsymbol{\theta}, \mathbf{y} \rangle) - \Omega(\boldsymbol{\mu}) = \langle \boldsymbol{\theta}', \boldsymbol{\mu} \rangle - \Omega(\boldsymbol{\mu})$.

We now proceed to the proof of Proposition 8. Suppose that L_Ω has the separation margin property, and let $\boldsymbol{\theta} = m\mathbf{y}$. Then, for any $\mathbf{y}' \in \mathcal{Y}$, we have $\langle \boldsymbol{\theta}, \mathbf{y}' \rangle + \frac{m}{2} \|\mathbf{y} - \mathbf{y}'\|^2 = m\langle \mathbf{y}, \mathbf{y}' \rangle + \frac{m}{2} \|\mathbf{y} - \mathbf{y}'\|^2 = \frac{m}{2} \|\mathbf{y}\|^2 + \frac{m}{2} \|\mathbf{y}'\|^2 = mr^2 = \langle m\mathbf{y}, \mathbf{y} \rangle = \langle \boldsymbol{\theta}, \mathbf{y} \rangle$, which implies that $L_\Omega(\boldsymbol{\theta}, \mathbf{y}) = 0$. From the first part of Proposition 2, this implies $m\mathbf{y} \in \partial\Omega(\mathbf{y})$.

Conversely, let us assume that $m\mathbf{y} \in \partial\Omega(\mathbf{y})$. From Lemma 2, this implies that $\boldsymbol{\theta} \in \partial\Omega(\mathbf{y})$ for any $\boldsymbol{\theta}$ such that $\langle \boldsymbol{\theta} - m\mathbf{y}, \mathbf{y}' \rangle \leq \langle \boldsymbol{\theta} - m\mathbf{y}, \mathbf{y} \rangle$ for all $\mathbf{y}' \in \mathcal{Y}$. Therefore, we have $\langle \boldsymbol{\theta}, \mathbf{y} \rangle - mr^2 \geq \langle \boldsymbol{\theta}, \mathbf{y}' \rangle - m\langle \mathbf{y}, \mathbf{y}' \rangle$, that is, $\langle \boldsymbol{\theta}, \mathbf{y} \rangle = \langle \boldsymbol{\theta}, \mathbf{y}' \rangle + \frac{m}{2} \|\mathbf{y} - \mathbf{y}'\|^2$. That is, any such $\boldsymbol{\theta}$ satisfies $\boldsymbol{\theta} \in \partial\Omega(\mathbf{y})$. From Proposition 2, this is equivalent to $L_\Omega(\boldsymbol{\theta}; \mathbf{y}) = 0$.

Let us now determine the margin of L_Ω , i.e., the smallest m such that $m\mathbf{y} \in \partial\Omega(\mathbf{y})$. From Lemma 2, this is equivalent to $mr^2 - \Omega(\mathbf{y}) \geq \langle m\mathbf{y}, \boldsymbol{\mu} \rangle - \Omega(\boldsymbol{\mu})$ for any $\boldsymbol{\mu} \in \text{conv}(\mathcal{Y})$, i.e., $\frac{-\Omega(\boldsymbol{\mu}) + \Omega(\mathbf{y})}{r^2 - \langle \mathbf{y}, \boldsymbol{\mu} \rangle} \leq m$, which leads to the expression (28).

B.8 Proof of Proposition 9

Let $\Omega(\mathbf{p}) = \sum_{j=1}^d g(p_j) + I_{\Delta^d}(\mathbf{p})$, where $g: [0, 1] \rightarrow \mathbb{R}_+$ is a non-negative, strictly convex, differentiable function. Therefore, g' is strictly monotonic on $[0, 1]$, thus invertible. We show how computing $\nabla(\Omega)^*$ reduces to finding the root of a monotonic scalar function, for which efficient algorithms are available.

From strict convexity and the definition of the convex conjugate,

$$\nabla(\Omega)^*(\boldsymbol{\theta}) = \operatorname{argmax}_{\mathbf{p} \in \Delta^d} \langle \mathbf{p}, \boldsymbol{\theta} \rangle - \sum_j g(p_j).$$

The constrained optimization problem above has Lagrangian

$$\mathcal{L}(\mathbf{p}, \boldsymbol{\nu}, \tau) := \sum_{j=1}^d g(p_j) - \langle \boldsymbol{\theta} + \boldsymbol{\nu}, \mathbf{p} \rangle + \tau(\mathbf{1}^\top \mathbf{p} - 1).$$

A solution $(\mathbf{p}^*, \boldsymbol{\nu}^*, \tau^*)$ must satisfy the KKT conditions

$$\begin{cases} g'(p_j) - \theta_j - \nu_j + \tau = 0 & \forall j \in [d] \\ \langle \mathbf{p}, \boldsymbol{\nu} \rangle = 0 \\ \mathbf{p} \in \Delta^d, \boldsymbol{\nu} \geq 0. \end{cases} \quad (44)$$

Let us define

$$\tau_{\min} := \max(\boldsymbol{\theta}) - g'(1) \quad \text{and} \quad \tau_{\max} := \max(\boldsymbol{\theta}) - g'\left(\frac{1}{d}\right).$$

Since g is strictly convex, g' is increasing and so $\tau_{\min} < \tau_{\max}$. For any $\tau \in [\tau_{\min}, \tau_{\max}]$, we construct $\boldsymbol{\nu}$ as

$$\nu_j := \begin{cases} 0, & \theta_j - \tau \geq g'(0) \\ g'(0) - \theta_j + \tau, & \theta_j - \tau < g'(0) \end{cases}$$

By construction, $\nu_j \geq 0$, satisfying dual feasibility. Injecting $\boldsymbol{\nu}$ into (44) and combining the two cases, we obtain

$$g'(p_j) = \max\{\theta_j - \tau, g'(0)\}. \quad (45)$$

We show that i) the stationarity conditions have a unique solution given τ , and ii) $[\tau_{\min}, \tau_{\max}]$ forms a sign-changing bracketing interval, and thus contains τ^* , which can then be found by one-dimensional search. The solution verifies all KKT conditions, thus is globally optimal.

Solving the stationarity conditions. Since g is strictly convex, its derivative g' is continuous and strictly increasing, and is thus a one-to-one mapping between $[0, 1]$ and $[g'(0), g'(1)]$. Denote by $(g')^{-1}: [g'(0), g'(1)] \rightarrow [0, 1]$ its inverse. If $\theta_j - \tau \geq g'(0)$, we have

$$\begin{aligned} g'(0) \leq g'(p_j) &= \theta_j - \tau \leq \max(\boldsymbol{\theta}) - \tau_{\min} \\ &= \max(\boldsymbol{\theta}) - \max(\boldsymbol{\theta}) + g'(1) \\ &= g'(1). \end{aligned}$$

Otherwise, $g'(p_j) = g'(0)$. This verifies that the r.h.s. of (45) is always within the domain of $(g')^{-1}$. We can thus apply the inverse to both sides to solve for p_j , obtaining

$$p_j(\tau) = (g')^{-1}(\max\{\theta_j - \tau, g'(0)\}). \quad (46)$$

Strict convexity implies the optimal \mathbf{p}^* is unique; it can be seen that τ^* is also unique. Indeed, assume optimal τ_1^*, τ_2^* . Then, $\mathbf{p}(\tau_1^*) = \mathbf{p}(\tau_2^*)$, so $\max(\boldsymbol{\theta} - \tau_1^*, g'(0)) = \max(\boldsymbol{\theta} - \tau_2^*, g'(0))$. This implies either $\tau_1^* = \tau_2^*$, or $\boldsymbol{\theta} - \tau_{\{1,2\}}^* \leq g'(0)$, in which case $\mathbf{p} = \mathbf{0} \notin \Delta^d$, which is a contradiction.

Validating the bracketing interval. Consider the primal infeasibility function $\phi(\tau) := \langle \mathbf{p}(\tau), \mathbf{1} \rangle - 1$; $\mathbf{p}(\tau)$ is primal feasible iff $\phi(\tau) = 0$. We show that ϕ is decreasing on $[\tau_{\min}, \tau_{\max}]$, and that it has opposite signs at the two extremities. From the intermediate value theorem, the unique root τ^* must satisfy $\tau^* \in [\tau_{\min}, \tau_{\max}]$.

Since g' is increasing, so is $(g')^{-1}$. Therefore, for all j , $p_j(\tau)$ is decreasing, and so is the sum $\phi(\tau) = \sum_j p_j(\tau) - 1$. It remains to check the signs at the boundaries.

$$\begin{aligned} \sum_i p_i(\tau_{\max}) &= \sum_i (g')^{-1}(\max\{\theta_i - \max(\boldsymbol{\theta}) + g'(1/d), g'(0)\}) \\ &\leq d (g')^{-1}(\max\{g'(1/d), g'(0)\}) \\ &= d (g')^{-1}(g'(1/d)) = 1, \end{aligned}$$

where we upper-bounded each term of the sum by the largest one. At the other end,

$$\begin{aligned} \sum_i p_i(\tau_{\min}) &= \sum_i (g')^{-1}(\max\{\theta_i - \max(\boldsymbol{\theta}) + g'(1), g'(0)\}) \\ &\geq (g')^{-1}(\max\{g'(1), g'(0)\}) \\ &= (g')^{-1}(g'(1)) = 1, \end{aligned}$$

using that a sum of non-negative terms is no less than its largest term. Therefore, $\phi(\tau_{\min}) \geq 0$ and $\phi(\tau_{\max}) \leq 0$. This implies that there must exist τ^* in $[\tau_{\min}, \tau_{\max}]$ satisfying $\phi(\tau^*) = 0$. The corresponding triplet $(\mathbf{p}(\tau^*), \boldsymbol{\nu}(\tau^*), \tau^*)$ thus satisfies all of the KKT conditions, confirming that it is the global solution.

Algorithm 1 is an example of a bisection algorithm for finding an approximate solution; more advanced root finding methods can also be used. We note that the resulting algorithm resembles the method provided in Krichene et al. (2015), with a non-trivial difference being the order of the thresholding and $(-g)^{-1}$ in Eq. (46).

Algorithm 1 Bisection for $\widehat{\mathbf{y}}_{\Omega}(\boldsymbol{\theta}) = \nabla \Omega^*(\boldsymbol{\theta})$

Input: $\boldsymbol{\theta} \in \mathbb{R}^d$, $\Omega(\mathbf{p}) = I_{\Delta^d} + \sum_i g(p_i)$

$\mathbf{p}(\tau) := (g')^{-1}(\max\{\boldsymbol{\theta} - \tau, g'(0)\})$

$\phi(\tau) := \langle \mathbf{p}(\tau), \mathbf{1} \rangle - 1$

$\tau_{\min} \leftarrow \max(\boldsymbol{\theta}) - g'(1)$;

$\tau_{\max} \leftarrow \max(\boldsymbol{\theta}) - g'(1/d)$

$\tau \leftarrow (\tau_{\min} + \tau_{\max})/2$

while $|\phi(\tau)| > \epsilon$

if $\phi(\tau) < 0$ $\tau_{\max} \leftarrow \tau$

else $\tau_{\min} \leftarrow \tau$

$\tau \leftarrow (\tau_{\min} + \tau_{\max})/2$

Output: $\nabla \widehat{\mathbf{y}}_{\Omega}(\boldsymbol{\theta}) \approx \mathbf{p}(\tau)$

B.9 Proof of Proposition 10

We first need the following lemma.

Lemma 3 *Let $\ell(\boldsymbol{\theta}; \mathbf{e}_k)$ be defined as*

$$\ell(\boldsymbol{\theta}; \mathbf{e}_k) := \begin{cases} \sum_j c_{k,j} \phi(\theta_j) & \text{if } \boldsymbol{\theta}^\top \mathbf{1} = 0 \\ \infty & \text{o.w.} \end{cases},$$

where $\phi: \mathbb{R} \rightarrow \mathbb{R}_+$ is convex. Then, H_ℓ defined in (34) equals

$$-H_\ell(\mathbf{p}) = \min_{\tau \in \mathbb{R}} \sum_j (\mathbf{p}^\top \mathbf{c}_j) \phi^* \left(\frac{-\tau}{\mathbf{p}^\top \mathbf{c}_j} \right).$$

Proof We want to solve

$$H_\ell(\mathbf{p}) = \min_{\boldsymbol{\theta}^\top \mathbf{1} = 0} \sum_j p_j \ell(\boldsymbol{\theta}; \mathbf{e}_j) = \min_{\boldsymbol{\theta}^\top \mathbf{1} = 0} \sum_j p_j \sum_i c_{j,i} \phi(\theta_i) = \min_{\boldsymbol{\theta}^\top \mathbf{1} = 0} \sum_i (\mathbf{p}^\top \mathbf{c}_i) \phi(\theta_i),$$

where \mathbf{c}_i is a vector gathering $c_{j,i}$ for all j . Introducing a Lagrange multiplier we get

$$H_\ell(\mathbf{p}) = \min_{\boldsymbol{\theta} \in \mathbb{R}^{|\mathcal{Y}|}} \max_{\tau \in \mathbb{R}} \sum_i (\mathbf{p}^\top \mathbf{c}_i) \phi(\theta_i) + \tau \sum_i \theta_i.$$

Strong duality holds and we can swap the order of the min and max. After routine calculations, we obtain

$$-H_\ell(\mathbf{p}) = \min_{\tau \in \mathbb{R}} \sum_i (\mathbf{p}^\top \mathbf{c}_i) \phi^* \left(\frac{-\tau}{\mathbf{p}^\top \mathbf{c}_i} \right). \quad (47)$$

■

We now prove Proposition 10. First, we rewrite $\ell(\boldsymbol{\theta}; \mathbf{e}_k) = \sum_{j \neq k} \phi(\theta_j - \theta_k)$ as $\ell(\boldsymbol{\theta}; \mathbf{e}_k) = \sum_j c_{k,j} \phi(\theta_j - \theta_k)$, where we choose $c_{k,j} = 0$ if $k = j$, 1 otherwise, leading to $\mathbf{p}^\top \mathbf{c}_j = 1 - p_j$. Because $\phi(\theta_j - \theta_k)$ is shift invariant w.r.t. $\boldsymbol{\theta}$, without loss of generality, we can further rewrite the loss as $\ell(\boldsymbol{\theta}; \mathbf{e}_k) = \sum_j c_{k,j} \phi(\theta_j)$ with $\text{dom}(\ell) = \{\boldsymbol{\theta} \in \mathbb{R}^{|\mathcal{Y}|} : \boldsymbol{\theta}^\top \mathbf{1} = 0\}$. Hence, Lemma 3 applies. We now derive closed form for specific choices of ϕ .

Hinge loss. When $\phi(t) = [1 + t]_+$, the conjugate is

$$\phi^*(u) = \begin{cases} -u & \text{if } u \in [0, 1] \\ \infty & \text{o.w.} \end{cases}.$$

The constraint set for τ is therefore $\mathcal{C} := \bigcap_{j \in [d]} [-\mathbf{p}^\top \mathbf{c}_j, 0] = \left[-\min_{j \in [d]} \mathbf{p}^\top \mathbf{c}_j, 0 \right]$. Hence

$$-H_\ell(\mathbf{p}) = \min_{\tau \in \mathcal{C}} d\tau = -d \min_{j \in [d]} \mathbf{p}^\top \mathbf{c}_j.$$

This recovers Duchi et al. (2018, Example 5, §A.6) with a simpler proof. We next turn to the following new results.

Smoothed hinge loss. We add quadratic regularization to the conjugate (Shalev-Shwartz and Zhang, 2016):

$$\phi^*(u) = \begin{cases} -u + \frac{1}{2}u^2 & \text{if } u \in [0, 1] \\ \infty & \text{o.w.} \end{cases}.$$

Going back to ϕ , we obtain:

$$\phi(t) = \begin{cases} 0 & \text{if } t \leq -1 \\ t + \frac{1}{2} & \text{if } t \geq 0 \\ \frac{1}{2}(1+t)^2 & \text{o.w.} \end{cases}. \quad (48)$$

The constraint set for τ is the same as before, $\mathcal{C} = \left[-\min_{j \in [d]} \mathbf{p}^\top \mathbf{c}_j, 0 \right]$.

Plugging ϕ^* into (47), we obtain

$$-\mathbf{H}_\ell(\mathbf{p}) = \min_{\tau \in \mathcal{C}} \frac{\tau^2}{2} \sum_{j=1}^d \frac{1}{\mathbf{p}^\top \mathbf{c}_j} + d \tau. \quad (49)$$

Since the problem is unidimensional, let us solve for τ unconstrained first:

$$\tau = -d / \left(\sum_{j=1}^d 1 / (\mathbf{p}^\top \mathbf{c}_j) \right). \quad (50)$$

We notice that $\tau \leq -\min_j \mathbf{p}^\top \mathbf{c}_j$ for $\mathbf{c}_j \geq \mathbf{0}$ since $\sum_j \frac{\min_i \mathbf{p}^\top \mathbf{c}_i}{\mathbf{p}^\top \mathbf{c}_j} \in [0, d]$. This expression of τ is not feasible. Hence the optimal solution is at the boundary and $\tau^* = -\min_j \mathbf{p}^\top \mathbf{c}_j$. Plugging that expression back into (49) gives the claimed expression of \mathbf{H}_ℓ .

Squared hinge loss. When $\phi(t) = \frac{1}{2}[1+t]_+^2$, the conjugate is

$$\phi^*(u) = \begin{cases} -u + \frac{1}{2}u^2 & \text{if } u \geq 0 \\ \infty & \text{o.w.} \end{cases}.$$

The constraint is now $\tau \leq 0$. Hence, the optimal solution of the problem w.r.t. τ in (47) is now (50) for all $\mathbf{p} \in \Delta^d, \mathbf{c}_j \geq \mathbf{0}$. Simplifying, we get

$$\mathbf{H}_\ell(\mathbf{p}) = \frac{\frac{1}{2}d^2}{\sum_{j=1}^d 1 / (\mathbf{p}^\top \mathbf{c}_j)}.$$

References

- Ryan P Adams and Richard S Zemel. [Ranking via sinkhorn propagation](#). *arXiv e-prints*, 2011.
- Nir Ailon, Kohei Hatano, and Eiji Takimoto. [Bandit online optimization over the permutahedron](#). *Theoretical Computer Science*, 650:92–108, 2016.
- Shun-ichi Amari. *Information Geometry and Its Applications*. Springer, 2016.
- Ehsan Amid and Manfred K Warmuth. [Two-temperature logistic regression based on the Tsallis divergence](#). *arXiv preprint arXiv:1705.07210*, 2017.
- Francis Bach, Simon Lacoste-Julien, and Guillaume Obozinski. [On the equivalence between herding and conditional gradient algorithms](#). *arXiv preprint arXiv:1203.4523*, 2012.
- Keith Ball, Eric A Carlen, and Elliott H Lieb. [Sharp uniform convexity and smoothness inequalities for trace norms](#). *Inventiones Mathematicae*, 115(1):463–482, 1994.
- Arindam Banerjee, Srujana Merugu, Inderjit S Dhillon, and Joydeep Ghosh. [Clustering with Bregman divergences](#). *JMLR*, 6:1705–1749, 2005.
- Ole Barndorff-Nielsen. *Information and Exponential Families: In Statistical Theory*. John Wiley & Sons, 1978.
- Leonard E. Baum and Ted Petrie. [Statistical inference for probabilistic functions of finite state Markov chains](#). *The Annals of Mathematical Statistics*, 37(6):1554–1563, 1966.
- Heinz H Bauschke and Patrick L Combettes. *Convex Analysis and Monotone Operator Theory in Hilbert Spaces*. Springer, 2nd edition, 2017.
- Amir Beck and Marc Teboulle. [A fast iterative shrinkage-thresholding algorithm for linear inverse problems](#). *SIAM Journal on Imaging Sciences*, 2(1):183–202, 2009.
- Amir Beck and Marc Teboulle. [Smoothing and first order methods: A unified framework](#). *SIAM Journal on Optimization*, 22(2):557–580, 2012.
- David Belanger, Dan Sheldon, and Andrew McCallum. [Marginal inference in MRFs using Frank-Wolfe](#). In *NeurIPS Workshop on Greedy Opt., FW and Friends*, 2013.
- Wolfgang H Berger and Frances L Parker. [Diversity of planktonic foraminifera in deep-sea sediments](#). *Science*, 168(3937):1345–1347, 1970.
- Dimitri P Bertsekas. *Nonlinear Programming*. Athena Scientific Belmont, 1999.
- Garrett Birkhoff. [Tres observaciones sobre el algebra lineal](#). *Univ. Nac. Tucumán Rev. Ser. A*, 5:147–151, 1946.
- Mathieu Blondel, Vivien Seguy, and Antoine Rolet. [Smooth and sparse optimal transport](#). In *Proc. of AISTATS*, 2018.

- Mathieu Blondel, André FT Martins, and Vlad Niculae. [Learning classifiers with Fenchel-Young losses: Generalized entropies, margins, and algorithms](#). In *Proc. of AISTATS*, 2019.
- Jonathan Borwein and Adrian S Lewis. *Convex Analysis and Nonlinear Optimization: Theory and Examples*. Springer Science & Business Media, 2010.
- Stephen Boyd and Lieven Vandenberghhe. *Convex Optimization*. Cambridge University Press, 2004.
- Lev M Bregman. [The relaxation method of finding the common point of convex sets and its application to the solution of problems in convex programming](#). *USSR Computational Mathematics and Mathematical Physics*, 7(3):200–217, 1967.
- Glenn W Brier. [Verification of forecasts expressed in terms of probability](#). *Monthly Weather Review*, 78(1):1–3, 1950.
- Peter Brucker. [An \$O\(n\)\$ algorithm for quadratic knapsack problems](#). *Operations Research Letters*, 3(3):163–166, 1984.
- Lars Buitinck, Gilles Louppe, Mathieu Blondel, Fabian Pedregosa, Andreas Mueller, Olivier Grisel, Vlad Niculae, Peter Prettenhofer, Alexandre Gramfort, Jaques Grobler, Robert Layton, Jake VanderPlas, Arnaud Joly, Brian Holt, and Gaël Varoquaux. [API design for machine learning software: experiences from the scikit-learn project](#). In *ECML PKDD Workshop: Languages for Data Mining and Machine Learning*, 2013.
- Andreas Buja, Werner Stuetzle, and Yi Shen. [Loss functions for binary class probability estimation and classification: Structure and applications](#). Technical report, University of Pennsylvania, 2005.
- Yoeng-Jin Chu and Tseng-Hong Liu. [On the shortest arborescence of a directed graph](#). *Science Sinica*, 14:1396–1400, 1965.
- Michael Collins. [Discriminative training methods for Hidden Markov Models: Theory and experiments with perceptron algorithms](#). In *Proc. of EMNLP*, 2002.
- Ronan Collobert, Jason Weston, Léon Bottou, Michael Karlen, Koray Kavukcuoglu, and Pavel Kuksa. [Natural language processing \(almost\) from scratch](#). *JMLR*, 12:2493–2537, 2011.
- Laurent Condat. [Fast projection onto the simplex and the \$\ell_1\$ ball](#). *Mathematical Programming*, 158(1-2):575–585, 2016.
- Koby Crammer and Yoram Singer. [On the algorithmic implementation of multiclass kernel-based vector machines](#). *JMLR*, 2:265–292, 2001.
- Imre Csiszár. [I-divergence geometry of probability distributions and minimization problems](#). *The Annals of Probability*, pages 146–158, 1975.
- Marco Cuturi. [Sinkhorn distances: Lightspeed computation of optimal transportation distances](#). In *Proc. of NeurIPS*, 2013.

- Marco Cuturi and Mathieu Blondel. [Soft-DTW: A differentiable loss function for time-series](#). In *Proc. of ICML*, 2017.
- John M Danskin. [The theory of max-min, with applications](#). *SIAM Journal on Applied Mathematics*, 14(4):641–664, 1966.
- George B Dantzig, Alex Orden, and Philip Wolfe. [The generalized simplex method for minimizing a linear form under linear inequality restraints](#). *Pacific Journal of Mathematics*, 5(2):183–195, 1955.
- Aaron Defazio, Francis Bach, and Simon Lacoste-Julien. [SAGA: A fast incremental gradient method with support for non-strongly convex composite objectives](#). In *Proc. of NeurIPS*, 2014.
- Morris H DeGroot. [Uncertainty, information, and sequential experiments](#). *The Annals of Mathematical Statistics*, pages 404–419, 1962.
- Arnaud Dessein, Nicolas Papadakis, and Jean-Luc Rouas. [Regularized optimal transport and the Rot Mover’s Distance](#). *arXiv preprint arXiv:1610.06447*, 2016.
- John C Duchi, Shai Shalev-Shwartz, Yoram Singer, and Tushar Chandra. [Efficient projections onto the \$\ell_1\$ -ball for learning in high dimensions](#). In *Proc. of ICML*, 2008.
- John C Duchi, Khashayar Khosravi, and Feng Ruan. [Multiclass classification, information, divergence, and surrogate risk](#). *The Annals of Statistics*, 46(6B):3246–3275, 2018.
- Joseph C Dunn and S Harshbarger. [Conditional gradient algorithms with open loop step size rules](#). *Journal of Mathematical Analysis and Applications*, 62(2):432–444, 1978.
- Jack Edmonds. [Optimum branchings](#). *J. Res. Nat. Bur. Stand.*, 71B:233–240, 1967.
- Jason Eisner. [Inside-outside and forward-backward algorithms are just backprop \(tutorial paper\)](#). In *Proc. of the Workshop on Structured Prediction for NLP*, 2016.
- Marguerite Frank and Philip Wolfe. [An algorithm for quadratic programming](#). *Naval Research Logistics Quarterly*, 3(1-2):95–110, 1956.
- Rafael Frongillo and Mark D Reid. [Convex foundations for generalized MaxEnt models](#). In *Proc. of AIP*, 2014.
- Dario Garcia-Garcia and Robert C Williamson. [Divergences and risks for multiclass experiments](#). In *Proc. of COLT*, 2012.
- Damien Garreau, Rémi Lajugie, Sylvain Arlot, and Francis Bach. [Metric learning for temporal sequence alignment](#). In *Proc. of NeurIPS*, 2014.
- Murray Gell-Mann and Constantino Tsallis. *Nonequilibrium Entropy: Interdisciplinary Applications*. Oxford University Press, 2004.
- Kevin Gimpel and Noah A Smith. [Softmax-margin CRFs: Training log-linear models with cost functions](#). In *Proc. of NAACL*, 2010.

- Corrado Gini. *Variabilità e mutabilità*. Reprinted in *Memorie di metodologica statistica* (Ed. Pizetti E, Salvemini, T). Rome: Libreria Eredi Virgilio Veschi, 1912.
- Tilmann Gneiting and Adrian E Raftery. Strictly proper scoring rules, prediction, and estimation. *Journal of the American Statistical Association*, 102(477):359–378, 2007.
- Arthur Gretton, Karsten M Borgwardt, Malte J Rasch, Bernhard Schölkopf, and Alexander Smola. A kernel two-sample test. *JMLR*, 13(Mar):723–773, 2012.
- Peter D Grünwald and A Philip Dawid. Game theory, maximum entropy, minimum discrepancy and robust Bayesian decision theory. *Annals of Statistics*, pages 1367–1433, 2004.
- Yann Guermeur. VC theory of large margin multi-category classifiers. *JMLR*, 8:2551–2594, 2007.
- David P Helmbold and Manfred K Warmuth. Learning permutations with exponential weights. *JMLR*, 10(Jul):1705–1736, 2009.
- Peter J Huber. Robust estimation of a location parameter. *The Annals of Mathematical Statistics*, 35(1):73–101, 1964.
- Martin Jaggi. Revisiting Frank-Wolfe: Projection-free sparse convex optimization. In *Proc. of ICML*, 2013.
- Thorsten Joachims. Optimizing search engines using clickthrough data. In *Proc. of ACM SIGKDD*. ACM, 2002.
- Roy Jonker and Anton Volgenant. A shortest augmenting path algorithm for dense and sparse linear assignment problems. *Computing*, 38(4):325–340, 1987.
- Dan Jurafsky and James H Martin. *Speech and Language Processing (3rd ed.)*. draft, 2018.
- Sham M Kakade, Karthik Sridharan, and Ambuj Tewari. On the complexity of linear prediction: Risk bounds, margin bounds, and regularization. In *Proc. of NeurIPS*. 2009.
- Diederik Kingma and Jimmy Ba. Adam: A method for stochastic optimization. In *ICLR*, 2015.
- Eliyahu Kiperwasser and Yoav Goldberg. Simple and accurate dependency parsing using bidirectional LSTM feature representations. *TACL*, 4:313–327, 2016.
- Gustav Kirchhoff. Ueber die auflösung der gleichungen, auf welche man bei der untersuchung der linearen vertheilung galvanischer ströme geführt wird. *Annalen der Physik*, 148(12): 497–508, 1847.
- Terry Koo, Amir Globerson, Xavier Carreras Pérez, and Michael Collins. Structured prediction models via the matrix-tree theorem. In *EMNLP*, 2007.
- Walid Krichene, Syrine Krichene, and Alexandre Bayen. Efficient Bregman projections onto the simplex. In *Proc. of CDC*. IEEE, 2015.

- Rahul G Krishnan, Simon Lacoste-Julien, and David Sontag. [Barrier Frank-Wolfe for marginal inference](#). In *Proc. of NeurIPS*, 2015.
- Harold W Kuhn. [The Hungarian method for the assignment problem](#). *Nav. Res. Log.*, 2(1-2):83–97, 1955.
- Anastasios Kyrillidis, Stephen Becker, Volkan Cevher, and Christoph Koch. [Sparse projections onto the simplex](#). In *Proc. of ICML*, 2013.
- Simon Lacoste-Julien and Martin Jaggi. [On the global linear convergence of Frank-Wolfe optimization variants](#). In *Proc. of NeurIPS*, 2015.
- Simon Lacoste-Julien, Martin Jaggi, Mark Schmidt, and Patrick Pletscher. [Block-coordinate Frank-Wolfe optimization for structural SVMs](#). In *Proc. of ICML*, 2012.
- Simon Lacoste-Julien, Fredrik Lindsten, and Francis Bach. [Sequential kernel herding: Frank-Wolfe optimization for particle filtering](#). In *Proc. of AISTATS*, 2015.
- John D Lafferty, Andrew McCallum, and Fernando CN Pereira. [Conditional Random Fields: Probabilistic models for segmenting and labeling sequence data](#). In *Proc. of ICML*, 2001.
- Guillaume Lample, Miguel Ballesteros, Sandeep Subramanian, Kazuya Kawakami, and Chris Dyer. [Neural architectures for named entity recognition](#). In *Proc. of NAACL*, 2016.
- Maksim Lapin, Matthias Hein, and Bernt Schiele. [Top-k multiclass SVM](#). In *Proc. of NeurIPS*, 2015.
- Cong Han Lim and Stephen J Wright. [Efficient bregman projections onto the permutahedron and related polytopes](#). In *Artificial Intelligence and Statistics*, pages 1205–1213, 2016.
- Dong C Liu and Jorge Nocedal. [On the limited memory BFGS method for large scale optimization](#). *Mathematical Programming*, 45(1):503–528, 1989.
- Francesco Locatello, Michael Tschannen, Gunnar Rätsch, and Martin Jaggi. [Greedy algorithms for cone constrained optimization with convergence guarantees](#). In *Proc. of NeurIPS*, 2017.
- Olvi L Mangasarian. [Pseudo-convex functions](#). *Journal of the Society for Industrial and Applied Mathematics, Series A: Control*, 3(2):281–290, 1965.
- André FT Martins and Ramón Fernandez Astudillo. [From softmax to sparsemax: A sparse model of attention and multi-label classification](#). In *Proc. of ICML*, 2016.
- André FT Martins, Mário AT Figueiredo, Pedro MQ Aguiar, Noah A Smith, and Eric P Xing. [Nonextensive information theoretic kernels on measures](#). *JMLR*, 10:935–975, 2009a.
- André FT Martins, Noah A Smith, and Eric P Xing. [Concise integer linear programming formulations for dependency parsing](#). In *ACL-IJCNLP*, 2009b.

- Hamed Masnadi-Shirazi. *The design of bayes consistent loss functions for classification*. PhD thesis, UC San Diego, 2011.
- Peter McCullagh and John A Nelder. *Generalized Linear Models*, volume 37. CRC press, 1989.
- Ryan T McDonald, Fernando CN Pereira, Kiril Ribarov, and Jan Hajič. Non-projective dependency parsing using spanning tree algorithms. In *Proc. of HLT-EMNLP*, 2005.
- Arthur Mensch and Mathieu Blondel. Differentiable dynamic programming for structured prediction and attention. In *Proc. of ICML*, 2018.
- Jean-Jacques Moreau. Proximité et dualité dans un espace hilbertien. *Bullet de la Société Mathématique de France*, 93(2):273–299, 1965.
- Renato Negrinho and Andre Martins. Orbit regularization. In *Proc. of NeurIPS*, 2014.
- John Ashworth Nelder and R Jacob Baker. *Generalized Linear Models*. Wiley Online Library, 1972.
- Yurii Nesterov. Smooth minimization of non-smooth functions. *Mathematical Programming*, 103(1):127–152, 2005.
- XuanLong Nguyen, Martin J Wainwright, and Michael I Jordan. On surrogate loss functions and f -divergences. *The Annals of Statistics*, 37(2):876–904, 2009.
- Vlad Niculae and Mathieu Blondel. A regularized framework for sparse and structured neural attention. In *Proc. of NeurIPS*, 2017.
- Vlad Niculae, André FT Martins, Mathieu Blondel, and Claire Cardie. SparseMAP: Differentiable sparse structured inference. In *Proc. of ICML*, 2018.
- Joakim Nivre, Marie-Catherine de Marneffe, Filip Ginter, Yoav Goldberg, Jan Hajic, Christopher D Manning, Ryan T McDonald, Slav Petrov, Sampo Pyysalo, Natalia Silveira, et al. Universal Dependencies v1: A multilingual treebank collection. In *LREC*, 2016.
- Jorge Nocedal and Stephen Wright. *Numerical Optimization*. Springer New York, 1999.
- Richard Nock and Frank Nielsen. Bregman divergences and surrogates for learning. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 31(11):2048–2059, 2009.
- Fabian Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot, and E. Duchesnay. Scikit-learn: Machine Learning in Python. *JMLR*, 12:2825–2830, 2011.
- Gabriel Peyré and Marco Cuturi. *Computational Optimal Transport*. 2017.
- Venkata Krishna Pillutla, Vincent Roulet, Sham M Kakade, and Zaid Harchaoui. A Smoother Way to Train Structured Prediction Models. In *Proc. of NeurIPS*, 2018.

- Mark D Reid and Robert C Williamson. [Composite binary losses](#). *JMLR*, 11:2387–2422, 2010.
- Mark D Reid and Robert C Williamson. [Information, divergence and risk for binary experiments](#). *JMLR*, 12:731–817, 2011.
- Alfréd Rényi. [On measures of entropy and information](#). In *Proc. of the 4th Berkeley Symposium on Mathematics, Statistics, and Probability*, volume 1, Berkeley, 1961. University of California Press.
- Ryan M Rifkin and Aldebaro Klautau. [In defense of one-vs-all classification](#). *JMLR*, 5(Jan):101–141, 2004.
- Ryan M Rifkin and Ross A Lippert. [Value regularization and Fenchel duality](#). *JMLR*, 8(Mar):441–479, 2007.
- R Tyrrell Rockafellar. *Convex Analysis*. Princeton University Press, 1970.
- Frank Rosenblatt. [The perceptron: a probabilistic model for information storage and organization in the brain](#). *Psychological Review*, 65(6):386, 1958.
- Hiroaki Sakoe and Seibi Chiba. [Dynamic programming algorithm optimization for spoken word recognition](#). *IEEE Transactions on Acoustics, Speech, and Signal Processing*, 26:43–49, 1978.
- Leonard J Savage. [Elicitation of personal probabilities and expectations](#). *Journal of the American Statistical Association*, 66(336):783–801, 1971.
- Bernhard Schölkopf and Alexander J Smola. *Learning With Kernels*. The MIT Press, Cambridge, MA, 2002.
- Shai Shalev-Shwartz and Sham M Kakade. [Mind the duality gap: Logarithmic regret algorithms for online optimization](#). In *Proc. of NeurIPS*, 2009.
- Shai Shalev-Shwartz and Yoram Singer. [Convex repeated games and Fenchel duality](#). In *Proc. of NeurIPS*, 2007.
- Shai Shalev-Shwartz and Yoram Singer. [On the equivalence of weak learnability and linear separability: New relaxations and efficient boosting algorithms](#). *Machine Learning*, 80(2-3):141–163, 2010.
- Shai Shalev-Shwartz and Tong Zhang. [Accelerated proximal stochastic dual coordinate ascent for regularized loss minimization](#). *Mathematical Programming*, 155(1):105–145, 2016.
- Claude E Shannon and Warren Weaver. *The Mathematical Theory of Communication*. University of Illinois Press, Urbana, Illinois, 1949.
- Chunhua Shen and Hanxi Li. [On the dual formulation of boosting algorithms](#). *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 32(12):2216–2231, 2010.

- Richard Sinkhorn and Paul Knopp. Concerning nonnegative matrices and doubly stochastic matrices. *Pacific Journal of Mathematics*, 21(2):343–348, 1967.
- David A Smith and Noah A Smith. Probabilistic models of nonprojective dependency trees. In *EMNLP*, 2007.
- Milan Straka and Jana Straková. Tokenizing, POS tagging, lemmatizing and parsing UD 2.0 with UDPipe. In *CoNLL Shared Task*, 2017.
- Daiki Suehiro, Kohei Hatano, Shuji Kijima, Eiji Takimoto, and Kiyohito Nagano. Online prediction under submodular constraints. In *International Conference on Algorithmic Learning Theory*. Springer, 2012.
- Hiroki Suyari. Generalization of Shannon-Khinchin axioms to nonextensive systems and the uniqueness theorem for the nonextensive entropy. *IEEE Trans. Information Theory*, 50(8):1783–1787, 2004.
- Robert E Tarjan. Finding optimum branchings. *Networks*, 7(1):25–35, 1977.
- Ben Taskar. *Learning Structured Prediction Models: A Large Margin Approach*. PhD thesis, Stanford University, 2004.
- Matus Telgarsky. A primal-dual convergence analysis of boosting. *JMLR*, 13(Mar):561–606, 2012.
- Constantino Tsallis. Possible generalization of Boltzmann-Gibbs statistics. *Journal of Statistical Physics*, 52:479–487, 1988.
- Ioannis Tsochantaridis, Thorsten Joachims, Thomas Hofmann, and Yasemin Altun. Large margin methods for structured and interdependent output variables. *JMLR*, 6:1453–1484, 2005.
- Nicolas Usunier, David Buffoni, and Patrick Gallinari. Ranking with ordered weighted pairwise classification. In *Proc. of ICML*, 2009.
- Leslie G Valiant. The complexity of computing the permanent. *Theor. Comput. Sci.*, 8(2): 189–201, 1979.
- Vladimir Vapnik. *Statistical Learning Theory*. Wiley, 1998.
- Marina Vinyes and Guillaume Obozinski. Fast column generation for atomic norm regularization, 2017.
- Andrew Viterbi. Error bounds for convolutional codes and an asymptotically optimum decoding algorithm. *IEEE Transactions on Information Theory*, 13(2):260–269, 1967.
- Martin J Wainwright and Michael I Jordan. Graphical models, exponential families, and variational inference. *Foundations and Trends® in Machine Learning*, 1(1–2):1–305, 2008.

- Manfred K Warmuth and Dima Kuzmin. [Randomized online PCA algorithms with regret bounds that are logarithmic in the dimension.](#) *JMLR*, 9:2287–2320, 2008.
- Robert C Williamson, Elodie Vernet, and Mark D Reid. [Composite multiclass losses.](#) *JMLR*, 2016.
- Philip Wolfe. [Finding the nearest point in a polytope.](#) *Mathematical Programming*, 11(1):128–149, 1976.
- Stephen J Wright, Robert D Nowak, and Mário AT Figueiredo. [Sparse reconstruction by separable approximation.](#) *IEEE Transactions on Signal Processing*, 57(7):2479–2493, 2009.
- Ronald R Yager. [On ordered weighted averaging aggregation operators in multicriteria decisionmaking.](#) *IEEE Transactions on Systems, Man, and Cybernetics*, 18(1):183–190, 1988.
- Shota Yasutake, Kohei Hatano, Shuji Kijima, Eiji Takimoto, and Masayuki Takeda. [Online linear optimization over permutations.](#) In *International Symposium on Algorithms and Computation*, pages 534–543. Springer, 2011.
- Daniel Zeman, Martin Popel, Milan Straka, Jan Hajic, Joakim Nivre, Filip Ginter, Juhani Luotolahti, Sampo Pyysalo, Slav Petrov, Martin Potthast, et al. [CoNLL 2017 shared task: Multilingual parsing from raw text to universal dependencies.](#) *CoNLL*, 2017.
- Xiangrong Zeng and Mário AT Figueiredo. [The ordered weighted \$\ell_1\$ norm: Atomic formulation and conditional gradient algorithm.](#) In *Proc. of SPARS*, 2015.
- Tong Zhang. [Statistical behavior and consistency of classification methods based on convex risk minimization.](#) *Annals of Statistics*, pages 56–85, 2004.