
Loss Functions and Operators Generated by f -Divergences

Vincent Roulet¹ Tianlin Liu¹ Nino Vieillard¹ Michaël E. Sander¹ Mathieu Blondel¹

Abstract

The logistic loss (a.k.a. cross-entropy loss) is one of the most popular loss functions used for multiclass classification. It is also the loss function of choice for next-token prediction in language modeling. It is associated with the Kullback–Leibler (KL) divergence and the softargmax operator. In this work, we build upon Fenchel–Young losses to construct convex loss functions generated from f -divergences. Our loss functions generalize the logistic loss in two directions: i) by replacing the KL divergence with f -divergences and ii) by allowing non-uniform reference measures. We instantiate our framework for numerous f -divergences, recovering existing losses and creating new ones. By analogy with the logistic loss, the loss function generated by an f -divergence is associated with an operator, that we dub f -softargmax. We derive a novel parallelizable bisection algorithm for computing the f -softargmax associated with any f -divergence. On the empirical side, one of the goals of this paper is to determine the effectiveness of loss functions beyond the classical cross-entropy in a language model setting, including on pre-training, post-training (SFT) and distillation. We show that the loss function generated by the α -divergence (which is equivalent to Tsallis α -negentropy in the case of unit reference measures) with $\alpha = 1.5$ performs well across several tasks.

1. Introduction

The logistic loss, a.k.a. cross-entropy loss, is widely used for multiclass classification. It is associated with the softargmax (a.k.a. softmax) operator, which turns logits into class probabilities. The logistic loss and the softargmax operators are also frequently used in the space of tokens for language model pre-training, post-training and distillation.

¹Google DeepMind. Correspondence to: Mathieu Blondel, Vincent Roulet <mblondel@google.com, vroulet@google.com>.

The classical softargmax is known to optimize a trade-off between the expected value of the logits and the Kullback–Leibler (KL) divergence with a uniform measure. This perspective is also adopted in the space of sequences, replacing the uniform measure with a reference measure, in reinforcement learning from human feedback (RLHF) (Christiano et al., 2017) and in direct preference optimization (DPO) (Rafailov et al., 2024). Recent works extended this perspective to general f -divergences (Go et al., 2023; Wang et al., 2024). However, these works naively compose an f -divergence with the (classical) softargmax, which does not result in a convex loss function.

In this work, building upon Fenchel–Young losses (Blondel et al., 2020), we propose to construct new convex losses based on f -divergences. Our losses generalize the logistic loss in two directions: i) by replacing the KL divergence with f -divergences and ii) by allowing non-uniform reference measures (class prior probabilities). Our loss construction generalizes the sparsemax (Martins & Astudillo, 2016) and entmax (Peters et al., 2019) losses, and allows us to create entirely new losses; see Table 1. In addition, each loss generated by an f -divergence is associated with a new operator for turning logits into probabilities, that we dub f -softargmax.

On the empirical side, one of the goals of this paper is to determine the effectiveness of loss functions beyond the classical cross-entropy in a language model setting, including on pre-training, post-training (SFT) and distillation.

To summarize, we make the following contributions.

- We propose to use f -divergence regularization to generate loss functions. Each loss function is associated with a new operator, that we dub f -softargmax. We instantiate this framework for numerous f -divergences, recovering existing losses and creating new ones.
- We derive a novel bisection algorithm for computing the f -softargmax associated with any f -divergence. Our algorithm parallelizes well on modern hardware.
- We demonstrate our loss functions on image classification, language model post-training and distillation.

Notation. Throughout this paper, k denotes the number of classes in a classification problem or the vocab-

ulary size in language modeling. We denote the set $[k] := \{1, \dots, k\}$. We denote the probability simplex by $\Delta^k := \{\mathbf{p} \in \mathbb{R}_+^k : \langle \mathbf{p}, \mathbf{1} \rangle = 1\}$. We denote the domain of f by $\text{dom}(f) := \{u : -\infty < f(u) < \infty\}$. We denote the convex conjugate of f by f^* , where $f^*(v) := \sup_u uv - f(u)$. We denote by $\boldsymbol{\theta} := h_{\mathbf{w}}(\mathbf{x})$ the logits produced by a network h with parameters $\mathbf{w} \in \mathcal{W}$ for the input $\mathbf{x} \in \mathcal{X}$. We denote hard (one-hot) labels as $\mathbf{y} \in \{\mathbf{e}_1, \dots, \mathbf{e}_k\}$, which are used for classification. More generally, we denote soft labels as $\mathbf{y} \in \Delta^k$, which are useful for learning from label proportions, as in distillation.

2. Background

2.1. Logistic loss

Given logits $\boldsymbol{\theta} \in \mathbb{R}^k$, we define the **softargmax** as

$$[\text{softargmax}(\boldsymbol{\theta})]_j := \frac{\exp(\theta_j)}{\sum_{j'=1}^k \exp(\theta_{j'})} \quad j \in [k]. \quad (1)$$

Note that (1) is commonly known in the literature as “softmax,” which is a misnomer as it is a smooth approximation of the argmax function (Blondel & Roulet, 2024). We instead define the **softmax** as

$$\text{softmax}(\boldsymbol{\theta}) := \log \sum_{j=1}^k \exp(\theta_j),$$

which is a smooth approximation of the maximum function. The softargmax and softmax are related by

$$\nabla \text{softmax} = \text{softargmax}.$$

The logistic loss (a.k.a. cross-entropy loss) between logits $\boldsymbol{\theta} \in \mathbb{R}^k$ and a ground-truth $\mathbf{y} \in \Delta^k$ is often defined as

$$-\sum_{j=1}^k y_j \log([\text{softargmax}(\boldsymbol{\theta})]_j) = \text{softmax}(\boldsymbol{\theta}) - \langle \boldsymbol{\theta}, \mathbf{y} \rangle.$$

The logistic loss is equivalent up to a constant to the KL divergence between the target and the softargmax output,

$$\begin{aligned} \ell(\boldsymbol{\theta}, \mathbf{y}) &:= \text{KL}(\mathbf{y}, \text{softargmax}(\boldsymbol{\theta})) \\ &= \text{softmax}(\boldsymbol{\theta}) - \langle \boldsymbol{\theta}, \mathbf{y} \rangle + \langle \mathbf{y}, \log \mathbf{y} \rangle. \end{aligned} \quad (2)$$

Adding the constant term $\langle \mathbf{y}, \log \mathbf{y} \rangle$ ensures that the loss is non-negative even when using soft labels (the constant term is zero with hard labels). The loss ℓ is convex in $\boldsymbol{\theta}$, and its gradient w.r.t. $\boldsymbol{\theta} \in \mathbb{R}^k$ is

$$\nabla_{\boldsymbol{\theta}} \ell(\boldsymbol{\theta}, \mathbf{y}) = \text{softargmax}(\boldsymbol{\theta}) - \mathbf{y}. \quad (3)$$

2.2. Binary logistic loss

In the special case $k = 2$, letting $\boldsymbol{\theta} := (0, \theta)$ and $\mathbf{y} := (1 - y, y)$ for $\theta \in \mathbb{R}$ and $y \in [0, 1]$, we obtain the **softplus**

$$\text{softplus}(\theta) := \log(1 + \exp(\theta)) = \text{softmax}(\boldsymbol{\theta})$$

and the **sigmoid**

$$\text{sigmoid}(\theta) := \frac{1}{1 + \exp(-\theta)} = [\text{softargmax}(\boldsymbol{\theta})]_1.$$

The two operators are again related by

$$\text{softplus}' = \text{sigmoid}.$$

The binary logistic loss between the logit $\theta \in \mathbb{R}$ and the ground-truth label $y \in [0, 1]$ is then

$$\ell(\theta, y) := \text{softplus}(\theta) + (1 - y) \log(1 - y) + y \log y - \theta y. \quad (4)$$

Beyond binary classification, a binary logistic loss can be used for pairwise ranking if we define $\theta := \theta_i - \theta_j$, where θ_i and θ_j are the scores of items $i \in [k]$ and $j \in [k]$, respectively. The probability that item $i \in [k]$ is ranked higher than item $j \in [k]$ according to the model is then

$$\text{sigmoid}(\theta_i - \theta_j) = \frac{\exp(\theta_i)}{\exp(\theta_i) + \exp(\theta_j)}.$$

This is called the Bradley-Terry model (1952).

2.3. Fenchel–Young losses

It is well-known that the softmax and softargmax can be written from a variational perspective as

$$\text{softmax}(\boldsymbol{\theta}) = \max_{\mathbf{p} \in \Delta^k} \langle \mathbf{p}, \boldsymbol{\theta} \rangle - \langle \mathbf{p}, \log \mathbf{p} \rangle$$

$$\text{softargmax}(\boldsymbol{\theta}) = \text{argmax}_{\mathbf{p} \in \Delta^k} \langle \mathbf{p}, \boldsymbol{\theta} \rangle - \langle \mathbf{p}, \log \mathbf{p} \rangle.$$

This suggests that we can create more general softmax and softargmax operators if we replace Shannon’s negative entropy $\langle \mathbf{p}, \log \mathbf{p} \rangle$ by more general regularization $\Omega: \Delta^k \rightarrow \mathbb{R}$, namely,

$$\text{softmax}_{\Omega}(\boldsymbol{\theta}) := \max_{\mathbf{p} \in \Delta^k} \langle \mathbf{p}, \boldsymbol{\theta} \rangle - \Omega(\mathbf{p})$$

$$\text{softargmax}_{\Omega}(\boldsymbol{\theta}) := \text{argmax}_{\mathbf{p} \in \Delta^k} \langle \mathbf{p}, \boldsymbol{\theta} \rangle - \Omega(\mathbf{p}).$$

By analogy with (2), we define the **Fenchel–Young** loss (Blondel et al., 2020) generated by Ω as

$$\ell_{\Omega}(\boldsymbol{\theta}, \mathbf{y}) := \text{softmax}_{\Omega}(\boldsymbol{\theta}) - \langle \boldsymbol{\theta}, \mathbf{y} \rangle + \Omega(\mathbf{y}) \quad (5)$$

$$\geq B_{\Omega}(\mathbf{y}, \text{softargmax}_{\Omega}(\boldsymbol{\theta})), \quad (6)$$

where B_{Ω} is the Bregman divergence generated by Ω . The loss in (5) is always convex w.r.t. $\boldsymbol{\theta}$, unlike the loss in (6). Among many other desirable properties, if Ω is strictly convex, Fenchel–Young losses satisfy

$$\ell_{\Omega}(\boldsymbol{\theta}, \mathbf{y}) = 0 \iff \text{softargmax}_{\Omega}(\boldsymbol{\theta}) = \mathbf{y} \quad (7)$$

and the gradient w.r.t. $\boldsymbol{\theta} \in \mathbb{R}^k$ is

$$\nabla_{\boldsymbol{\theta}} \ell_{\Omega}(\boldsymbol{\theta}, \mathbf{y}) = \text{softargmax}_{\Omega}(\boldsymbol{\theta}) - \mathbf{y}.$$

In the binary classification setting, we can similarly replace sigmoid with sigmoid_{Ω} in (4), as well as the regularization term $(1 - y) \log(1 - y) + y \log y$ with $\Omega((1 - y, y))$.

Table 1. By using f -divergences as regularization, we can generalize existing loss functions (logistic, sparsemax, entmax) to non-uniform reference measure \mathbf{q} and we can construct several new loss functions. Reverse f -divergences (not listed below), such as the reverse KL, can also be used to generate alternative loss functions. The f -softargmax operators associated with the Chi-square divergence and the α -divergence for $\alpha > 1$ can produce probability distributions with sparse support. When using the unit positive measure $\mathbf{q} = \mathbf{1}$, we obtain f -entropies, which recover existing known entropies (Shannon, Gini, Tsallis) up to a constant.

Divergence	Entropy	Loss	Sparse
Kullback–Leibler	Shannon	Logistic	No
Chi-square	Gini	Sparsemax	Yes
α -divergence	Tsallis	Entmax	$\alpha > 1$
Jensen–Shannon	–	–	No
Squared Hellinger	–	–	No

3. Generating losses from f -divergences

In this paper, we propose to study Fenchel–Young losses and associated operators when the regularizer is defined as

$$\Omega_f(\mathbf{p}; \mathbf{q}) := D_f(\mathbf{p}, \mathbf{q}), \quad (8)$$

where D_f is a f -divergence and $\mathbf{q} \in \mathbb{R}_+^k$ is a reference measure, which contains the prior class weights. When \mathbf{q} is not available, we can simply use $\mathbf{q} = \mathbf{1}$, in which case we will obtain negative f -entropies, as explained below.

3.1. f -divergences

Let $f: \mathbb{R}_+ \rightarrow \mathbb{R}$ be a convex function such that $f(1) = 0$ and $f(0) = \lim_{u \rightarrow 0^+} f(u)$. The f -divergence between two discrete positive measures $\mathbf{p} \in \mathbb{R}_+^k$ and $\mathbf{q} \in \mathbb{R}_+^k$ (Rényi, 1961; Csiszár, 1967; Ali & Silvey, 1966) is then

$$D_f(\mathbf{p}, \mathbf{q}) := \sum_{j=1}^k f(p_j/q_j) q_j = \langle f(\mathbf{p}/\mathbf{q}), \mathbf{q} \rangle,$$

where f and division are applied element-wise. An f -divergence is always non-negative and jointly convex in \mathbf{p} and \mathbf{q} . Many existing divergences can be written in f -divergence form: see Table 1 and Appendix B.

Divergence reversal. Since an f -divergence is not necessarily symmetric in \mathbf{p} and \mathbf{q} , it may seem that setting the reference measure \mathbf{q} as right argument in (8) is arbitrary. However, if we define $g(u) := uf(1/u)$, we obtain the reverse divergence of D_f , $D_g(\mathbf{p}, \mathbf{q}) = D_f(\mathbf{q}, \mathbf{p})$. Therefore, for any f -divergence, we can always construct the corresponding reverse divergence, without loss of generality.

3.2. f -entropies

As we emphasized, our proposed framework can take into account a discrete reference measure $\mathbf{q} \in \mathbb{R}_+^k$, which intuitively contains the prior class weights. When such a measure is not available, we can simply choose the unit positive measure $\mathbf{q} = \mathbf{1}$. In this case, we recover negative f -entropies (Cichocki & Amari, 2010), a.k.a. f -negentropies. They are defined as

$$\Omega_f(\mathbf{p}) := \sum_{j=1}^k f(p_j) = \langle f(\mathbf{p}), \mathbf{1} \rangle = D_f(\mathbf{p}, \mathbf{1}).$$

By choosing f , we recover (up to a constant) numerous existing negentropies. When $f(u) = u \log u$, which is the generating function of the KL divergence, we recover the **Shannon negentropy**,

$$\Omega_f(\mathbf{p}) = \langle \mathbf{p}, \log \mathbf{p} \rangle.$$

When $f(u) = \frac{1}{2}(u^2 - 1)$, which is the generating function of the Chi-square divergence, we recover the **Gini negentropy**, used for the sparsemax loss (Martins & Astudillo, 2016),

$$\Omega_f(\mathbf{p}) \doteq \frac{1}{2}(\|\mathbf{p}\|_2^2 - 1)$$

(we use \doteq for equality up to a constant). More generally, with $f(u) = \frac{(u^\alpha - 1) - \alpha(u - 1)}{\alpha(\alpha - 1)}$, which is the generating function of α -divergences (using Cichocki & Amari (2010)’s definition), we recover the **Tsallis negentropy**,

$$\Omega_f(\mathbf{p}) \doteq \frac{1}{\alpha} \langle \mathbf{p}, \log_\alpha(\mathbf{p}) \rangle = \frac{1}{\alpha(\alpha - 1)} (\|\mathbf{p}\|_\alpha^\alpha - 1),$$

used by the entmax loss (Peters et al., 2019). The Tsallis negentropy itself is very general, as it recovers the Shannon negentropy when $\alpha \rightarrow 1$ and the Gini negentropy when $\alpha = 2$. In our experiments, we will demonstrate good results with the choice $\alpha = 1.5$, which can be thought as a middle ground between Shannon entropy (used by the logistic loss) and Gini entropy (used by the sparsemax loss).

Unexplored entropy instances. Beyond existing entropies, some choices of f lead to entropies that had not been considered before, such as the squared Hellinger, illustrated in Figure 1. We leave a formal axiomatic analysis of such new entropies to future work, and refer the interested reader to, e.g., Ben-Bassat (1978); Friedland et al. (2013); Sharma & Sazim (2021); Sbert et al. (2024).

Effective domain. We point out, however, that some f -entropies are only well-defined on the relative interior of the probability simplex, if $\lim_{u \rightarrow 0} f(u) = -\infty$. This is for instance the case of the f -entropies associated with the reverse KL and Jeffrey divergences. This means that the loss functions generated by these choices only work with strictly positive soft labels.




Figure 1. Illustration of f -entropies $-\Omega_f(\mathbf{p})$ for $\mathbf{p} = (1 - \pi, \pi)$ and varying $\pi \in [0, 1]$. We add a constant $f(0)$ to ensure non-negativity of the f -entropies.

3.3. f -softmax and f -softargmax

Overloading the notation, we define the f -softmax as

$$\text{softmax}_f(\boldsymbol{\theta}; \mathbf{q}) := \max_{\mathbf{p} \in \Delta^k} \langle \mathbf{p}, \boldsymbol{\theta} \rangle - D_f(\mathbf{p}, \mathbf{q}) \in \mathbb{R} \quad (9)$$

and the f -softargmax as

$$\text{softargmax}_f(\boldsymbol{\theta}; \mathbf{q}) := \underset{\mathbf{p} \in \Delta^k}{\text{argmax}} \langle \mathbf{p}, \boldsymbol{\theta} \rangle - D_f(\mathbf{p}, \mathbf{q}) \in \Delta^k. \quad (10)$$

Compared to a classical softmax and softargmax, our operators use a function f and include an additional reference measure \mathbf{q} as argument. When $f(u) = u \log u$ and $\mathbf{q} = \mathbf{1}$, we recover the classical softmax and softargmax.

Scaling the divergence by a temperature parameter $\beta > 0$ can easily be done. Indeed, for any $\beta > 0$,

$$\begin{aligned} \text{softmax}_{\beta f}(\boldsymbol{\theta}; \mathbf{q}) &= \beta \text{softmax}_f(\boldsymbol{\theta}/\beta; \mathbf{q}) \\ \text{softargmax}_{\beta f}(\boldsymbol{\theta}; \mathbf{q}) &= \text{softargmax}_f(\boldsymbol{\theta}/\beta; \mathbf{q}). \end{aligned}$$

Sparse distributions. As summarized in Table 1, the f -softargmax associated with the Chi-square and α divergences for $\alpha > 1$ can produce probability distributions with sparse support, meaning that some classes have exactly zero probability according to the model. As will be clear from Proposition 1, the f -softargmax can be sparse when $0 \in \text{dom}(f') \iff \lim_{u \rightarrow 0} f'(u) > -\infty$.

3.4. f -softplus and f -sigmoid

As special cases of f -softmax and f -softargmax, by defining $\boldsymbol{\theta} := (0, \theta)$, $\mathbf{p} := (1 - \pi, \pi)$ and $\mathbf{q} := (q_0, q_1)$, we obtain the f -softplus

$$\text{softplus}_f(\theta; \mathbf{q}) := \max_{\pi \in [0, 1]} \pi \theta - D_f((1 - \pi, \pi), \mathbf{q}) \in \mathbb{R}$$

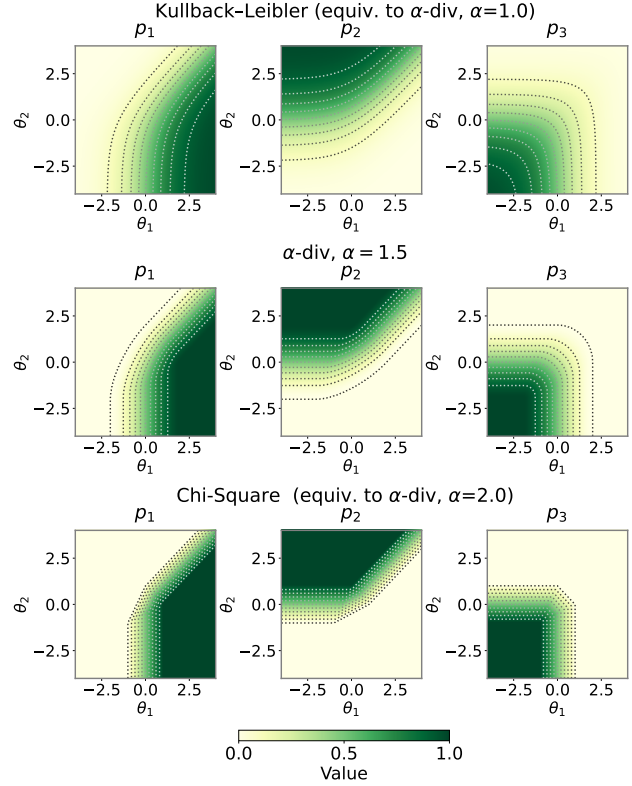


Figure 2. Illustration of $(p_1, p_2, p_3) = \text{softargmax}_f(\theta_1, \theta_2, 0; \mathbf{q})$ when varying $\theta_1, \theta_2 \in \mathbb{R}$ for three possible divergences and with $\mathbf{q} = (1, 1, 1)$. More illustrations are given in Appendix B.

and the f -sigmoid

$$\text{sigmoid}_f(\theta; \mathbf{q}) := \underset{\pi \in [0, 1]}{\text{argmax}} \pi \theta - D_f((1 - \pi, \pi), \mathbf{q}) \in [0, 1].$$

The effect of the prior class weights $\mathbf{q} \in \mathbb{R}_+^2$ on the shape of the sigmoid is illustrated in Figure 8 in Appendix B. The f -softplus and f -sigmoid can be used, not only for binary classification, but also for pairwise ranking, by analogy with Section 2.2. Similarly to the f -softmax and f -softargmax, we can easily scale D_f by a temperature parameter $\beta > 0$.

3.5. Loss function

To obtain a loss function associated with the f -softargmax, we instantiate the Fenchel–Young loss defined in (5) with the regularization defined in (8) to define

$$\begin{aligned} \ell_f(\boldsymbol{\theta}, \mathbf{y}; \mathbf{q}) &:= \text{softmax}_f(\boldsymbol{\theta}; \mathbf{q}) + \Omega_f(\mathbf{y}; \mathbf{q}) - \langle \boldsymbol{\theta}, \mathbf{y} \rangle \\ &= \text{softmax}_f(\boldsymbol{\theta}; \mathbf{q}) + D_f(\mathbf{y}, \mathbf{q}) - \langle \boldsymbol{\theta}, \mathbf{y} \rangle, \end{aligned} \quad (11)$$

where $\boldsymbol{\theta} \in \mathbb{R}^k$ are the logits, $\mathbf{y} \in \Delta^k$ is the ground-truth and $\mathbf{q} \in \mathbb{R}_+^k$ is a reference measure, which contains the prior class weights. This loss inherits from all the desirable properties of Fenchel–Young losses. In particular, it is convex w.r.t. $\boldsymbol{\theta}$ and it is differentiable everywhere if f

is strictly convex. On the other hand, it is not necessarily convex w.r.t. \mathbf{q} , since softmax_f involves the maximum over a collection of *concave* functions of \mathbf{q} . Because $\mathbf{y} \in \triangle^k$ and not just $\mathbf{y} \in \{e_1, \dots, e_k\}$, this loss can also be used for learning from label proportions, as is useful in distillation for example.

Choosing \mathbf{q} . One distinctive feature of our losses compared to usual instances of Fenchel–Young losses is the possibility to adjust the reference measure \mathbf{q} . In the absence of prior knowledge, we can simply use $\mathbf{q} = \mathbf{1}$ or $\mathbf{q} = \mathbf{1}/k$. This recovers Fenchel–Young losses generated by f -negentropies. If class weights are available, we can use this prior knowledge as \mathbf{q} .

3.6. Computation

On first sight, it is not obvious how to solve the variational problems (9) and (10) involved in computing the f -softmax and f -softargmax, respectively. We need an algorithm that works for any valid choice of f and parallelizes well, as we often need to compute the operators on a batch of b logits $\theta_{i_1}, \dots, \theta_{i_b}$, where $\theta_{i_j} := h_{\mathbf{w}}(\mathbf{x}_{i_j})$.

In this section, we introduce a novel generic algorithm for computing the f -softmax and f -softargmax given access to f and f^* , where f^* is the convex conjugate of f . The latter is usually available in closed form; we give numerous examples in Appendix B.

For convenience, with a slight abuse of notation, we define the shorthand $f'_* := (f^*)'$. We denote $f'(0) = \lim_{x \rightarrow 0, x \geq 0} f'(x)$ which may be finite or $-\infty$.

Proposition 1. Reduction to root finding

Let f be a strictly convex and differentiable function such that $(0, +\infty) \subseteq \text{dom } f'$. Then, for any $\boldsymbol{\theta} \in \mathbb{R}^k$,

$$\text{softmax}_f(\boldsymbol{\theta}; \mathbf{q}) = \tau^* + \sum_{j=1}^k q_j f'_*(\max\{\theta_j - \tau^*, f'(0)\}) \quad (12)$$

$$[\text{softargmax}_f(\boldsymbol{\theta}; \mathbf{q})]_j = q_j f'_*(\max\{\theta_j - \tau^*, f'(0)\}), \quad (13)$$

where τ^* is the unique solution of

$$\sum_{j=1}^k q_j f'_*(\max\{\theta_j - \tau, f'(0)\}) = 1, \quad (14)$$

on $\tau \in [\tau_{\min}, \tau_{\max}]$, where, for $j^* \in \arg\max_{j \in [k]} \theta_j$,

$$\begin{aligned} \tau_{\min} &:= \theta_{j^*} - f'(1/q_{j^*}) \\ \tau_{\max} &:= \theta_{j^*} - f'\left(1/\left(\sum_{j=1}^k q_j\right)\right). \end{aligned}$$

A proof is given in Appendix C.4. Proposition 1 is a generalization of (Blondel et al., 2020, Proposition 9) to arbitrary

reference measures \mathbf{q} . Our proof technique is different: it uses Fenchel duality as opposed to Lagrange duality and it rigorously accounts for the domain of f' , which is one of the technical difficulties for supporting general f -divergences. Proposition 1 is a generalization of (Wang et al., 2024, Theorem 1) to the case $0 \in \text{dom}(f')$, that is, to f -softargmax operators with sparse output. A Newton algorithm was proposed in (Terjék, 2021) for computing the same operator. However, that algorithm was used for regularized optimal transport, not for creating loss functions. In addition, it assumes that $(f + \delta_{\mathbb{R}_+})^*$ is twice differentiable, which is not the case when $0 \in \text{dom}(f')$, and does not enjoy convergence guarantees, as a line search was not used. See also the discussion in (Belousov & Peters, 2017) on the implications of $0 \notin \text{dom}(f')$.

Implementation. In practice, we solve the root equation in (14) by bisection (Algorithm 1). This algorithm has several advantages: it is simple, parallelizes well on GPU and TPU (we often need to compute the f -softargmax of a batch) and achieves an error on iteration t of $(\tau_{\max} - \tau_{\min})/2^t$. That is, the error exponentially decreases with the number of iterations. We show in Appendix A.4 that the overhead of our proposed algorithm is negligible compared to a classical softargmax.

Algorithm 1 Computing f -softmax and f -softargmax

Inputs: logits $\boldsymbol{\theta} \in \mathbb{R}^k$, prior $\mathbf{q} \in \mathbb{R}_+^k$, tolerance $\epsilon > 0$

$[\mathbf{p}(\tau)]_j := q_j f'_*(\max\{\theta_j - \tau, f'(0)\})$, $j \in [k]$

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

$j^* \in \arg\max_{j \in [k]} \theta_j$

$\tau_{\min} \leftarrow \theta_{j^*} - f'(1/q_{j^*})$

$\tau_{\max} \leftarrow \theta_{j^*} - f'\left(1/\left(\sum_{j=1}^k q_j\right)\right)$

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

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

if $\phi(\tau) < 0$

$\tau_{\max} \leftarrow \tau$

else

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

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

Outputs: $\text{softargmax}_f(\boldsymbol{\theta}; \mathbf{q}) \approx \mathbf{p}(\tau)$

$\text{softmax}_f(\boldsymbol{\theta}; \mathbf{q}) \approx \tau + \sum_{j=1}^k q_j f'_*(\max\{\theta_j - \tau, f'(0)\})$

Closed forms in the binary case. In the special case $k = 2$ (binary classification), we can often derive closed-form solutions for the f -softplus and the f -sigmoid operator. For completeness, we derive the expressions for numerous cases in Appendix C.3.

3.7. Differentiation

Differentiating through f -softmax and loss. In order to differentiate the f -softmax, we can simply use Danskin’s theorem to obtain

$$\begin{aligned}\nabla_{\theta} \text{softmax}_f(\theta; q) &= p^* := \text{softargmax}_f(\theta; q) \\ \nabla_q \text{softmax}_f(\theta; q) &= -\nabla_q D_f(p^*, q).\end{aligned}$$

As a result, the loss gradients are

$$\begin{aligned}\nabla_{\theta} \ell_f(\theta, y; q) &= p^* - y = \text{softargmax}_f(\theta; q) - y \\ \nabla_q \ell_f(\theta, y; q) &= \nabla_q D_f(y, q) - \nabla_q D_f(p^*, q).\end{aligned}$$

The first equation is in complete analogy with (3). The second equation is the difference of the f -divergence gradients evaluated at the ground-truth y and the prediction p^* . Differentiating through the f -softplus operator and its associated loss function is similar.

Differentiating through f -softargmax. When the goal is to use the f -softargmax as the operator associated with our loss functions, we do not need to differentiate through the f -softargmax. As explained above, thanks to Danskin’s theorem, differentiating through the f -softmax is sufficient. When the goal is to use the f -softargmax as an attention mechanism, however, we do need to differentiate through the f -softargmax. This is more challenging as, from Proposition 1, we need to differentiate through the solution of a root equation. Under assumptions on f , we can apply the implicit function theorem (Krantz & Parks, 2002) through the root equation’s solution. This can be implemented using automatic implicit differentiation (Blondel et al., 2022). Importantly, implicit differentiation does not require solving a costly linear system here. Since the root equation is one-dimensional, a simple division is sufficient.

4. Experiments

To evaluate different f -divergence generated losses, we apply them to tasks of different data modalities, including image classification (Section 4.1) and text generation (Section 4.2). These experiments also cover different training strategies: from scratch, finetuning, and distillation.

4.1. ImageNet classification

We apply different f -divergence generated losses to train a ResNet50 model (He et al., 2016) on the ImageNet-2012 dataset (Russakovsky et al., 2015). The ImageNet dataset contains 1.28 million training images and 50,000 validation images, belonging to one of 1,000 classes. The ResNet50 model is a standard choice for ImageNet.

We use an SGD optimizer with 0.9 momentum to train the ResNet50 model for 90 epochs. During the initial 5 epochs,

Table 2. ImageNet classification results.

Divergence	Accuracy (%)
Kullback–Leibler	76.87
Chi-square	76.06
α -divergence ($\alpha = 1.5$)	77.56
Jensen–Shannon	72.24
Squared Hellinger	72.81

we use a linear warmup that achieves a peak learning rate of 0.2; we then use cosine annealing to reduce the learning rate to 0. The weight decay is set to be 10^{-4} . The batch size is 512.

Table 2 shows the validation accuracy of training ResNet50 using different f -divergence generated losses; boldface indicates the highest accuracy. Recall that the KL generated loss is equivalent to the standard cross-entropy loss; therefore, using the KL generated loss should match the result of prior work (He et al., 2016). Our experiments confirm this: the KL loss reaches 76.87% accuracy, consistent with previous benchmarks (Appendix A.1). Perhaps surprisingly, we find that the α -divergence loss function surpasses the KL loss in validation accuracy. This improvement is achieved without adjusting any hyperparameters and therefore is notable. In addition, $\alpha = 1.5$ seems to be optimal within α -divergences as shown in Figure 5, where we see that validation accuracy is maximal for $\alpha = 1.5$ among 11 values for $\alpha \in [1, 2]$. However, the Jensen–Shannon, Squared Hellinger, and Chi-square (Pearson divergence) generated loss functions perform worse than the KL divergence in this experiment.

4.2. Language modeling

We compare the performance of different f -divergence-based loss functions for training language models (LMs). To provide a thorough evaluation, we consider three common LM training strategies: (i) pretraining, (ii) supervised fine-tuning (SFT), and (iii) distillation-based fine-tuning. All three approaches optimize a next-token prediction loss:

$$\mathbb{E}_{(x, y)} \left[\sum_{t=1}^{|y|} \ell_f(h_w(x, y_{<t}), \phi(y_t); q) \right], \quad (15)$$

where h_w is the LM that outputs logits with parameters w , $|y|$ is the length of the text sequence y , and the reference measure $q = \mathbf{1}$ is a unitary prior; cf. Equation (11). Pretraining and SFT use one-hot vectors $\phi(y_t) \in \{e_1, \dots, e_k\}$ as the targets; distillation uses soft labels $\phi(y_t) \in \Delta^k$ obtained as the next-token probability of a teacher model.

Pretraining with f -divergence generated losses. We used the same pretraining method as the NanoDO models

(Liu et al., 2024; Wortsman et al., 2024), a set of well-tuned decoder-only transformer models trained on the public C4 dataset (Raffel et al., 2020). Specifically, we follow the training set up of the 1.2B parameter model from Wortsman et al. (2024); we use 250B tokens of C4 to train the model, so that the model is Chinchilla optimal (Hoffmann et al., 2022). Our pretraining results (see Appendix A.2) echo our main findings from the ImageNet experiments in Section 4.1. Specifically, losses based on α -divergence, with $\alpha = 1.5$, performed best for predicting the next word, slightly outperforming the standard KL divergence loss.

Learning q . Since our loss functions take a reference distribution q and we can compute gradients with respect to q , a natural idea is to try to learn it from data. As a negative result, using the NanoDO implementation above on the C4 dataset, we found that learning q does not seem to help improve performance compared to $q = 1$, both when using the KL divergence and when using the $\alpha = 1.5$ divergence.

SFT and distillation with f -divergence generated losses.

Instead of pretraining LMs from scratch, practitioners often fine-tune readily available, open-weight models pretrained with standard cross-entropy. This prompts a question: Can we effectively fine-tune these cross-entropy pretrained models using different f -divergence generated losses? To explore this, we evaluated f -divergence generated losses with two common finetuning methods, SFT and distillation, on a text summarization task (Narayan et al., 2018). For SFT, we use a pretrained T5-base model (Raffel et al., 2020) with 250M parameters and train it on the XSum dataset (Narayan et al., 2018) with a next-token prediction loss (15). Traditionally, the cross-entropy loss is used as the next-token prediction loss, which is equivalent to KL-divergence. In our experiments, however, we evaluate more general f -divergence generated losses.

Distillation uses a large teacher model’s class probabilities as soft labels to train a smaller student model (Hinton et al., 2015). In the LM context, such soft labels are the teacher model’s next-token probabilities, which can be more effective than one-hot labels as they provide richer information about the likely next tokens. We use a T5-XL model (800M parameters) as the teacher and a T5-base model (250M parameters) as the student; prior to distillation, both models were SFT’ed on XSum as in Agarwal et al. (2024). We then fit the student model with the loss (15), where $\phi(y_t)$ in the loss are the soft labels produced by the teacher model.

For both SFT and distillation, we adopted the XSum settings from Agarwal et al. (2024, Appendix A.3). Following Chowdhery et al. (2023) and Agarwal et al. (2024), we evaluated summarization quality using the ROUGE-2 score (Lin, 2004) with temperature-based sampling (temperature=1). Figure 3 shows results for SFT (unhatched) and distilla-




Figure 3. Comparing the effect of different f -divergence generated losses used for SFT and distillation training of LMs. The f -divergence generated losses are Jensen-Shannon (js), reverse Chi-square (rcs), Kullback-Leibler (kl), Chi-square (cs), and α -divergence with $\alpha = 1.5$ (alpha1.5)

tion (hatched). Distillation consistently outperformed SFT. Critically, Chi-squared and α -divergence ($\alpha = 1.5$) losses outperformed KL divergence for both training paradigms.

Comparing f -softmax variants for decoding. In addition to the loss function used at the training of LMs, another key factor that influences the performance of LMs is how one decode responses from them. While the training of a LM uses a fixed f -divergence loss, when decoding, one can choose different f -softmax variants to turn LM logits into next-token probabilities for sampling. In our prior experiments of SFT and distillation (Figure 3), at decoding time we used the f -softmax corresponding to each f -divergence loss function used at training time. But this raises a question: Do the performance differences primarily come from the f -divergence loss used during training or from the usage of different f -softmax at decoding? To study this, we perform SFT with different f -divergence generated losses as before, but then decode from the trained LM using either the associated f -softmax or the classical (KL-based) softmax. Figure 4 shows the results. Decoding with the two types of f -softmax variants yields nearly identical results. Indeed, note that the two bars of KL divergence are slightly different only due to the effect of random sampling; the observed performance differences in the case of other divergences are comparable to those seen in the KL case. This suggests that the f -divergence used during training primarily contributes to the performance differences of models.




Figure 4. Decoding with the classical (KL-based) softmax performs similarly to decoding with f -softmax associated with the loss used for training. This suggests that the choice of f -divergence loss during training, not the decoding method, primarily drives performance differences.

4.3. Summary of empirical findings

We summarize the main take-aways from our experiments.

- We found that the loss function generated by the α -divergence with $\alpha = 1.5$ works well across all tasks we tried. While equivalent to Tsallis negentropy ($\alpha = 1.5$) with a uniform reference measure (Blondel et al., 2020; Peters et al., 2019), its effectiveness across several language modeling tasks is novel. Its good performance could intuitively come from the fact that it is a middle ground between the logistic loss ($\alpha = 1$) and the sparse-max loss ($\alpha = 2$). Figure 5 confirms that accuracy is maximized around $\alpha = 1.5$.
- We successfully fine-tuned LMs with various f -divergence generated losses, despite the fact the pretraining was carried out using the cross-entropy loss. This enables direct application of our f -divergence generated losses to pretrained, open-weight LMs.
- We show that, during text generation, standard softmax yielded good performance even when the model is fine-tuned with f -divergence generated losses. This is surprising because (7) tells us that we should in principle use the f -softmax operator associated with the loss. This empirical finding opens up the possibility to use our losses without changing the inference code, which is convenient when working with open-weight LMs.
- Loss functions generated by f -divergences not considered before (Jensen-Shannon, squared Hellinger) did not lead to better accuracy. The KL and $\alpha = 1.5$ generated losses consistently achieved the best results across tasks.

5. Related work

Loss functions based on f -divergences. Nguyen et al. (2009) studied surrogate losses for binary classification based on f -divergences. However, this approach is not straightforward to generalize to the multiclass setting, as it requires multi-distribution extensions of f -divergences (Garcia-Garcia & Williamson, 2012; Duchi et al., 2018). Recently, f -divergences have been used for distillation in combination with a classical softmax (Agarwal et al., 2024). In our notation, this defines a loss $\theta \mapsto D_f(y, \text{softmax}(\theta))$. However, this does not result in a convex loss in θ . Novello & Tonello (2024) developed deep learning objectives based on the variational formulation of f -divergences. Sbert et al. (2024); Sharma & Sazim (2021); Aminian et al. (2024) also defined entropies from f -divergence to measure uncertainty, homogeneity or to build losses robust to label errors as also done by Zhu et al. (2023). f -divergences have also been used to design GANs (Nowozin et al., 2016) or more generally provide alternate learning frameworks (Li & Turner, 2016).

Belousov (2017) briefly studied the idea of generating Bregman divergences from f -divergences. However, Bregman divergences need to be explicitly composed with a softmax and this composition is again usually not convex. In contrast, in our approach, each loss $\theta \mapsto \ell_f(\theta, y; q)$ is convex and is implicitly associated with the corresponding softmax $_f$ operator.

LM alignment with f -divergences. In RLHF, given a learned reward model $r: \mathcal{X} \times \mathcal{Y} \rightarrow \mathbb{R}$, the optimal policy is

$$p^*(\cdot|x) := \arg\max_{p \in \Delta^{|\mathcal{Y}|}} \mathbb{E}_{y \sim p}[r(x, y)] - \text{KL}(p, q(\cdot|x)),$$

where $x \in \mathcal{X}$ is a given prompt and $q(\cdot|x) \in \Delta^{|\mathcal{Y}|}$ is a reference conditional distribution, usually a pre-trained model. This approach can be generalized to f -divergences (Go et al., 2023). In our notation, this can be written as

$$p^*(\cdot|x) = \text{softmax}_f(r(x, \cdot), q(\cdot|x)).$$

To avoid learning a separate reward model altogether, we can use direct preference optimization (DPO) (Rafailov et al., 2024), which was generalized to f -divergences (Wang et al., 2024). Proposition 1 can be seen as generalization of (Wang et al., 2024, Theorem 1) that supports $0 \in \text{dom}(f')$. In this paper, our focus was on the SFT and distillation steps, which usually precede RLHF. Unlike these works, we propose to use the f -softmax operator and the corresponding Fenchel–Young loss for pretraining, SFT and distillation in the space of tokens, making SGD-based optimization easy.

6. Conclusion

We proposed to use f -divergences as regularization for generating loss functions and associated operators (f -softmax, f -softargmax, f -softplus and f -sigmoid). Our proposal establishes a link between the logistic, sparsemax and entmax loss functions and the KL, Chi-square and α divergences, respectively. Thanks to this perspective, our proposal generalizes these losses to non-uniform class weights. Other choices of f -divergences allowed us to create entirely new loss functions. Overall, we found that the loss function generated by the $\alpha = 1.5$ divergence worked comparably or better to the cross-entropy loss on all tasks we tried.

Impact statement

This paper explores f -divergence generated loss functions for classification or language modelling. We do not foresee any specific ethical or societal implications arising directly from this work.

Acknowledgments

We thank the anonymous reviewers for their comments that helped improve the paper.

References

- Agarwal, R., Vieillard, N., Zhou, Y., Stanczyk, P., Garea, S. R., Geist, M., and Bachem, O. On-policy distillation of language models: learning from self-generated mistakes. In *Proceedings of the International Conference on Learning Representations (ICLR)*, 2024.
- Ali, S. M. and Silvey, S. D. A general class of coefficients of divergence of one distribution from another. *Journal of the Royal Statistical Society: Series B (Methodological)*, 28(1):131–142, 1966.
- Aminian, G., Bagheri, A., JafariNodeh, M., Karimian, R., and Yassaee, M.-H. Robust semi-supervised learning via f -divergence and α -rényi divergence. In *Proceedings of the IEEE International Symposium on Information Theory (ISIT)*, pp. 1842–1847. IEEE, 2024.
- Belousov, B. Bregman divergence of alpha-divergence, 2017. URL <http://www.boris-belousov.net/2017/04/16/bregman-divergence/>.
- Belousov, B. and Peters, J. f -divergence constrained policy improvement. *arXiv preprint arXiv:1801.00056*, 2017.
- Ben-Bassat, M. f -entropies, probability of error, and feature selection. *Information and Control*, 39(3):227–242, 1978.
- Blondel, M. and Roulet, V. The elements of differentiable programming. *arXiv preprint arXiv:2403.14606*, 2024.
- Blondel, M., Martins, A. F., and Niculae, V. Learning with Fenchel–Young losses. *Journal of Machine Learning Research*, 21(35):1–69, 2020.
- Blondel, M., Berthet, Q., Cuturi, M., Frostig, R., Hoyer, S., Llinares-Lopez, F., Pedregosa, F., and Vert, J.-P. Efficient and modular implicit differentiation. In *Advances in Neural Information Processing Systems*, 2022.
- Bradley, R. A. and Terry, M. E. Rank analysis of incomplete block designs: I. the method of paired comparisons. *Biometrika*, 39(3/4):324–345, 1952.
- Chowdhery, A., Narang, S., Devlin, J., Bosma, M., Mishra, G., Roberts, A., Barham, P., Chung, H. W., Sutton, C., Gehrmann, S., et al. PaLM: Scaling language modeling with pathways. *Journal of Machine Learning Research*, 24(240):1–113, 2023.
- Christiano, P. F., Leike, J., Brown, T., Martic, M., Legg, S., and Amodei, D. Deep reinforcement learning from human preferences. In *Advances in Neural Information Processing Systems*, 2017.
- Cichocki, A. and Amari, S.-i. Families of alpha- beta- and gamma-divergences: Flexible and robust measures of similarities. *Entropy*, 12(6):1532–1568, 2010.
- Csiszár, I. On information-type measure of difference of probability distributions and indirect observations. *Studia Sci. Math. Hungar.*, 2:299–318, 1967.
- Dehghani, M., Djolonga, J., Mustafa, B., Padlewski, P., Heek, J., Gilmer, J., Steiner, A. P., Caron, M., Geirhos, R., Alabdulmohsin, I., Jenatton, R., Beyer, L., Tschannen, M., Arnab, A., Wang, X., Riquelme Ruiz, C., Minderer, M., Puigcerver, J., Evci, U., Kumar, M., Steenkiste, S. V., Elsayed, G. F., Mahendran, A., Yu, F., Oliver, A., Huot, F., Bastings, J., Collier, M., Gritsenko, A. A., Birodkar, V., Vasconcelos, C. N., Tay, Y., Mensink, T., Kolesnikov, A., Pavetic, F., Tran, D., Kipf, T., Lucic, M., Zhai, X., Keysers, D., Harmsen, J. J., and Houlsby, N. Scaling vision transformers to 22 billion parameters. In *Proceedings of the International Conference on Machine Learning (ICML)*, pp. 7480–7512, 2023.
- Duchi, J., Khosravi, K., and Ruan, F. Multiclass classification, information, divergence and surrogate risk. *The Annals of Statistics*, 46(6B):3246 – 3275, 2018.
- Friedland, S., Gheorghiu, V., and Gour, G. Universal uncertainty relations. *Physical review letters*, 111(23), 2013.
- Garcia-Garcia, D. and Williamson, R. C. Divergences and risks for multiclass experiments. In *Proceedings of the Annual Conference on Learning Theory (COLT)*, pp. 28–1. JMLR Workshop and Conference Proceedings, 2012.

- Go, D., Korbak, T., Kruszewski, G., Rozen, J., Ryu, N., and Dymetman, M. Aligning language models with preferences through f -divergence minimization. In *Proceedings of the International Conference on Machine Learning (ICML)*, 2023.
- He, K., Zhang, X., Ren, S., and Sun, J. Deep residual learning for image recognition. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 770–778, 2016.
- Hinton, G., Vinyals, O., and Dean, J. Distilling the knowledge in a neural network. *arXiv preprint arXiv:1503.02531*, 2015.
- Hoffmann, J., Borgeaud, S., Mensch, A., Buchatskaya, E., Cai, T., Rutherford, E., de Las Casas, D., Hendricks, L. A., Welbl, J., Clark, A., Hennigan, T., Noland, E., Millican, K., van den Driessche, G., Damoc, B., Guy, A., Osindero, S., Simonyan, K., Elsen, E., Vinyals, O., Rae, J., and Sifre, L. An empirical analysis of compute-optimal large language model training. In *Advances in Neural Information Processing Systems*, 2022.
- Krantz, S. G. and Parks, H. R. *The implicit function theorem: history, theory, and applications*. Springer Science & Business Media, 2002.
- Li, Y. and Turner, R. E. Rényi divergence variational inference. *Advances in neural information processing systems*, 2016.
- Lin, C.-Y. Rouge: A package for automatic evaluation of summaries. In *Text Summarization Branches Out*, pp. 74–81, 2004.
- Liu, P. J., Novak, R., Lee, J., Wortsman, M., Xiao, L., Everett, K., Alemi, A. A., Kurzeja, M., Marcenac, P., Gur, I., Kornblith, S., Xu, K., Elsayed, G., Fischer, I., Pennington, J., Adlam, B., and Dickstein, J.-S. Nanodo: A minimal transformer decoder-only language model implementation in JAX., 2024. URL <http://github.com/google-deepmind/nanodo>.
- Martins, A. and Astudillo, R. From softmax to sparsemax: A sparse model of attention and multi-label classification. In *Proceedings of the International Conference on Machine Learning (ICML)*, 2016.
- Narayan, S., Cohen, S. B., and Lapata, M. Don’t give me the details, just the summary! Topic-aware convolutional neural networks for extreme summarization. In *Proceedings of the Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pp. 1797–1807. Association for Computational Linguistics, 2018.
- Nguyen, X., Wainwright, M. J., and Jordan, M. I. On surrogate loss functions and f -divergences. *The Annals of Statistics*, 37(2):876 – 904, 2009.
- Novello, N. and Tonello, A. M. f -divergence based classification: beyond the use of cross-entropy. In *Proceedings of the International Conference on Machine Learning (ICML)*, 2024.
- Nowozin, S., Cseke, B., and Tomioka, R. f -gan: Training generative neural samplers using variational divergence minimization. *Advances in neural information processing systems*, 2016.
- Peters, B., Niculae, V., and Martins, A. F. Sparse sequence-to-sequence models. In *Proceedings of the Annual Meeting of the Association for Computational Linguistics (ACL)*, 2019.
- Radford, A., Wu, J., Child, R., Luan, D., Amodei, D., Sutskever, I., et al. Language models are unsupervised multitask learners. *OpenAI blog*, 1(8):9, 2019.
- Rafailov, R., Sharma, A., Mitchell, E., Manning, C. D., Ermon, S., and Finn, C. Direct preference optimization: Your language model is secretly a reward model. In *Advances in Neural Information Processing Systems*, 2024.
- Raffel, C., Shazeer, N., Roberts, A., Lee, K., Narang, S., Matena, M., Zhou, Y., Li, W., and Liu, P. J. Exploring the limits of transfer learning with a unified text-to-text transformer. *Journal of Machine Learning Research*, 21 (140):1–67, 2020.
- Rényi, A. On measures of entropy and information. In *Proceedings of the fourth Berkeley symposium on mathematical statistics and probability, volume 1: contributions to the theory of statistics, volume 4*, pp. 547–562. University of California Press, 1961.
- Russakovsky, O., Deng, J., Su, H., Krause, J., Satheesh, S., Ma, S., Huang, Z., Karpathy, A., Khosla, A., Bernstein, M., Berg, A. C., and Fei-Fei, L. ImageNet large scale visual recognition challenge. *International Journal of Computer Vision*, 115:211–252, 2015.
- Sbert, M., Chen, M., Poch, J., Feixas, M., and Chen, S. Entropies from f -divergences. In *Proceedings of the International Conference on Modeling Decisions for Artificial Intelligence (MDAI)*, pp. 47–60. Springer, 2024.
- Sharma, G. and Sazim, S. A geometric way to find the measures of uncertainty from statistical divergences for discrete and finite probability distributions. *arXiv preprint arXiv:2106.14874*, 2021.
- Su, J., Ahmed, M., Lu, Y., Pan, S., Bo, W., and Liu, Y. Roformer: Enhanced transformer with rotary position embedding. *Neurocomputing*, 568:127063, 2024.
- Terjék, D. Moreau–Yosida f -divergences. In *Proceedings of the International Conference on Machine Learning (ICML)*, 2021.

Wang, C., Jiang, Y., Yang, C., Liu, H., and Chen, Y. Beyond reverse KL: Generalizing direct preference optimization with diverse divergence constraints. In *Proceedings of the International Conference on Learning Representations (ICLR)*, 2024.

Wortsman, M., Liu, P. J., Xiao, L., Everett, K. E., Alemi, A. A., Adlam, B., Co-Reyes, J. D., Gur, I., Kumar, A., Novak, R., Pennington, J., Sohl-Dickstein, J., Xu, K., Lee, J., Gilmer, J., and Kornblith, S. Small-scale proxies for large-scale transformer training instabilities. In *Proceedings of the International Conference on Learning Representations (ICLR)*, 2024.

Zhu, D., Ying, Y., and Yang, T. Label distributionally robust losses for multi-class classification: Consistency, robustness and adaptivity. In *Proceedings of the International Conference on Machine Learning (ICML)*, 2023.

A. Experimental details and additional experiments

A.1. ImageNet classification

When using the KL divergence, our training yields 76.87% accuracy on ImageNet. This result matches previous benchmarks of ResNet50 trained on ImageNet, such as <https://github.com/google/flax/tree/main/examples/imagenet>.




Figure 5. Validation accuracy on ImageNet when using the α -divergence generated loss, with $\alpha \in [1., 2.]$. Accuracy is maximized near $\alpha = 1.5$.

A.2. Pretraining NanoDO models

The 1.2B NanoDO model’s architecture and hyperparameters follow from the setup in Wortsman et al. (2024). The model is similar to GPT2 (Radford et al., 2019) but incorporates modern features like rotary positional embeddings (Su et al., 2024) and qk-layernorm (Dehghani et al., 2023). To evaluate the trained model, we use the next-token prediction accuracy on the validation dataset. Our results (Table 3) show that α -divergence is on par, and slightly more performant, than the classical KL approach.

Table 3. Next-token prediction accuracy

Divergence	Accuracy (%)
Kullback–Leibler	48.66
Chi-square	46.75
α -divergence ($\alpha = 1.5$)	48.70

A.3. Standard deviation of accuracy results over random seeds

Table 4. Mean, standard deviation, min, max accuracies over 5 random seeds for the ImageNet classification results.

Divergence	Mean	Standard Deviation	Min	Max
Kullback–Leibler	0.7684	0.0007	0.7676	0.7692
Chi-square	0.7604	0.0015	0.7587	0.7621
α -divergence ($\alpha = 1.5$)	0.7758	0.0013	0.7743	0.7776
Jensen–Shannon	0.7246	0.0014	0.723	0.7266
Squared Hellinger	0.7281	0.0008	0.7272	0.729

Table 5. Mean, standard deviation, min, max Rouge-2 scores over 5 random seeds for the supervised fine-tuning (SFT) results.

Divergence	Mean	Standard Deviation	Min	Max
Kullback–Leibler	9.77	0.02	9.75	9.8
Chi–square	11.15	0.1	11.02	11.31
α -divergence ($\alpha = 1.5$)	14.27	0.04	14.2	14.32
Jensen–Shannon	7.95	0.08	7.87	8.07
Reverse Chi–squared	9.55	0.1	9.44	9.7

Table 6. Mean, standard deviation, min, max Rouge-2 scores over 5 random seeds for the distillation results.

Divergence	Mean	Standard Deviation	Min	Max
Kullback–Leibler	16.64	0.05	16.57	16.71
Chi–square	14.17	0.09	14.01	14.26
α -divergence ($\alpha = 1.5$)	17.43	0.13	17.19	17.6
Jensen–Shannon	16.51	0.06	16.43	16.6
Reverse Chi–squared	16.3	0.05	16.25	16.38

A.4. Computational cost

Bisection convergence. We experimentally validate the exponential convergence of our proposed Algorithm 1. Our results are in Figure 6 and confirm the theoretical convergence rate.




Figure 6. Error for computing the root in Proposition 1 using our bisection based Algorithm 1 as a function of the number of iterations. We use the output with 30 iterations as a proxy for the true root. We compare the measured error (dark purple) with the theoretical 2^{-t} error (light purple). This experiment is run on TPU.

Overall cost. We propose an experiment validating that our proposed losses generated by f -divergences and associated operators do not affect the runtime of standard learning pipelines. For this, we consider a Residual Network (He et al., 2016) ResNet18, with either the standard softmax or the softmax associated with the $\alpha = 1.5$ divergence. For different batch sizes, we profile the time needed for running both the model and the f -softmax. Results in Figure 7 show that the runtimes are approximately equal.




Figure 7. Profiled time for an input of shape $(b, 224, 224, 3)$ to be processed by a ResNet18 followed by f -softargmax corresponding the 1.5-divergence (dark purple) or the standard Kullback–Leibler divergence (light purple, in which case we use the standard JAX implementation) for different batches size b . This experiment is run on TPU.

A.5. Impact of prior distribution on sigmoids




Figure 8. Illustration of $\theta \mapsto \text{sigmoid}_f(\theta; (1 - q_1, q_1))$ for $q_1 \in \{0.1, 0.25, 0.5, 0.75, 0.9\}$ and for $f(u) = u \log u$, the generating function of the KL divergence. We see that $\text{sigmoid}_f(\theta; (1 - q_1, q_1))$ at $\theta = 0$ is equal to q_1 . Intuitively, if a model is uncertain and produces a value of $\theta = 0$, then the model outputs q_1 , the prior probability of the positive class.

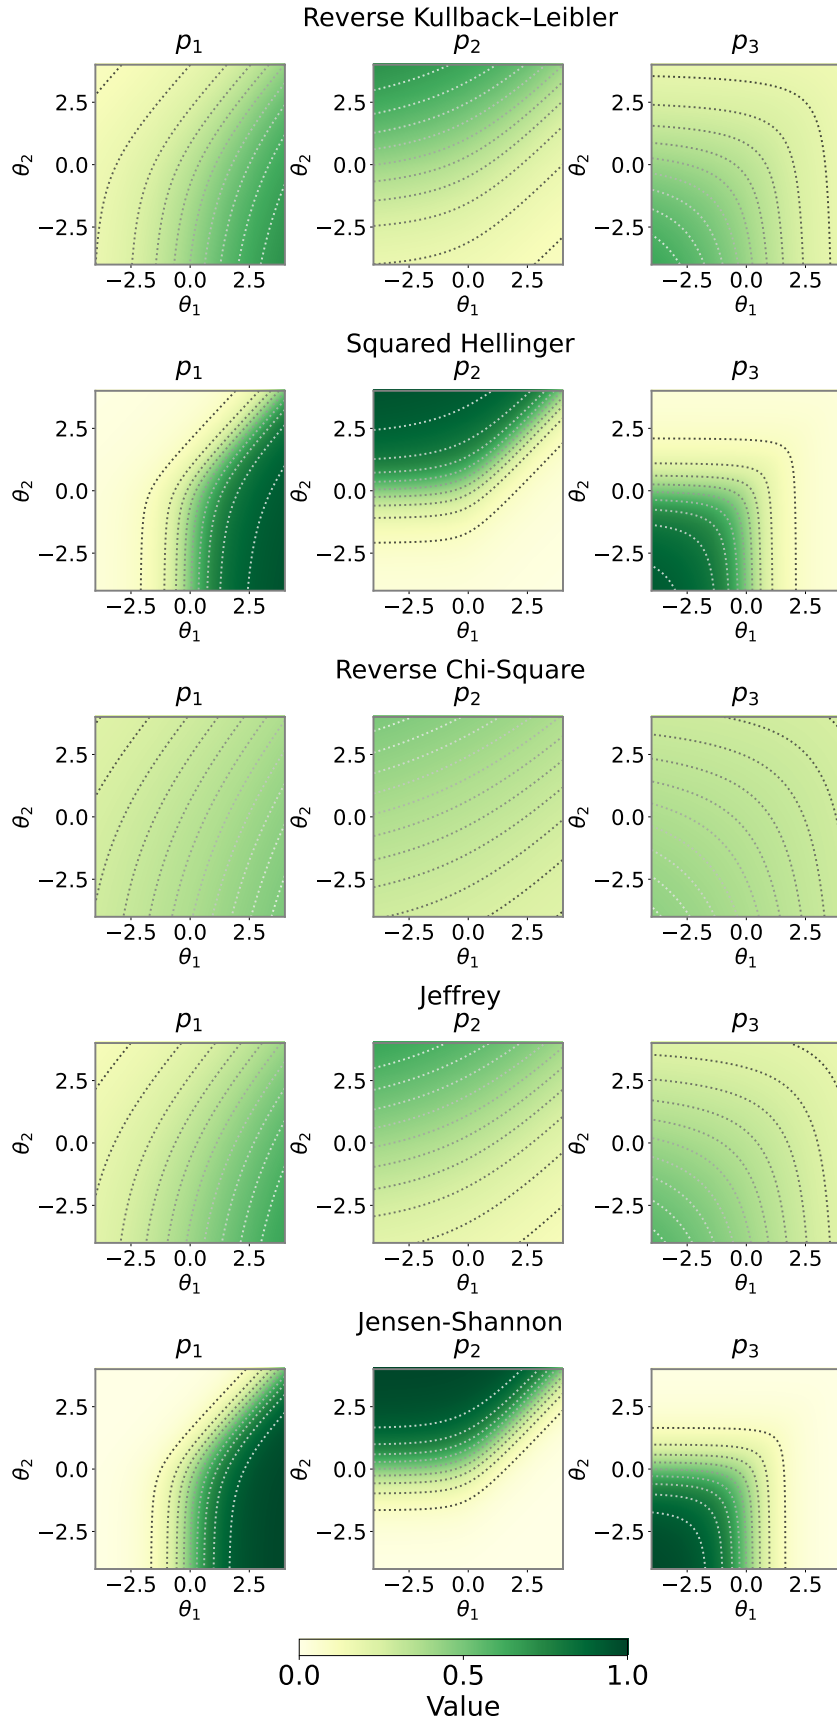


Figure 9. Illustration of $(p_1, p_2, p_3) = \text{softargmax}_f(\theta_1, \theta_2, 0; \mathbf{q})$ when varying $\theta_1, \theta_2 \in \mathbb{R}$ for diverse f -divergences and with $\mathbf{q} = (1, 1, 1)$.

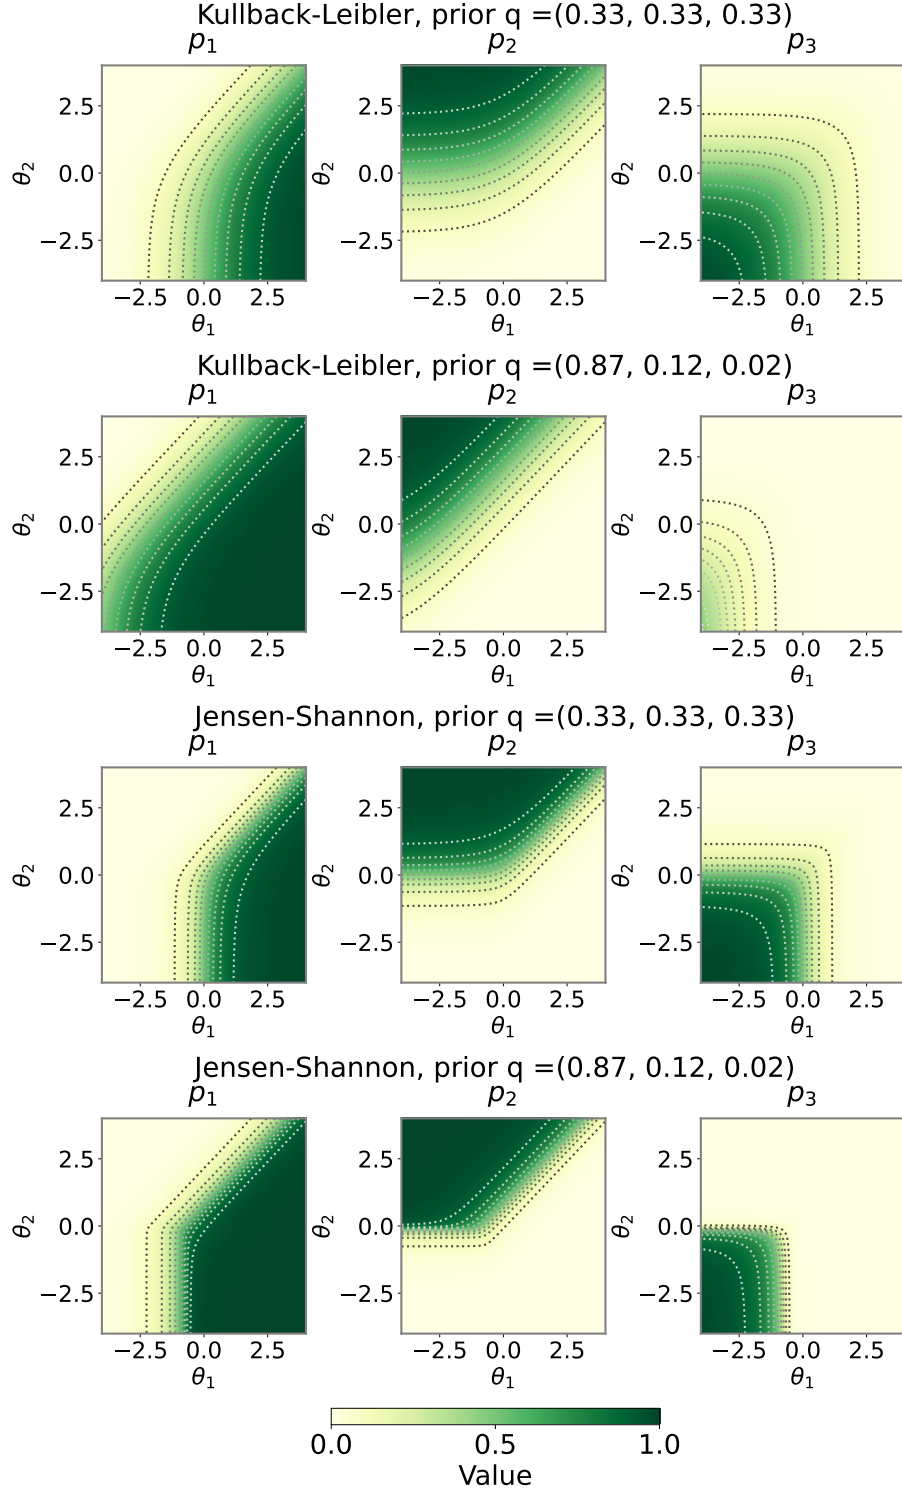


Figure 10. Illustration of $(p_1, p_2, p_3) = \text{softargmax}_f(\theta_1, \theta_2, 0; \mathbf{q})$ when varying $\theta_1, \theta_2 \in \mathbb{R}$ for diverse f -divergences and priors \mathbf{q} .

B. Examples of f -divergences

- KL

- $f(u) = u \log u$, $\text{dom}(f) = \mathbb{R}_+$
- $f'(u) = \log u + 1$
- $f_*(v) = \exp(v - 1)$, $\text{dom}(f_*) = \mathbb{R}$
- $f'_*(v) = \exp(v - 1)$
- $D_f(p, q) = \text{KL}(p, q) := \langle p, \log p \rangle - \langle p, \log q \rangle$

- Generalized KL

- $f(u) = u \log u - (u - 1)$, $\text{dom}(f) = \mathbb{R}_+$
- $f'(u) = \log u$
- $f_*(v) = \exp(v) - 1$, $\text{dom}(f_*) = \mathbb{R}$
- $f'_*(v) = \exp(v)$
- $D_f(p, q) = \text{GKL}(p, q) := \langle p, \log p \rangle - \langle p, \log q \rangle - \langle p, 1 \rangle + \langle q, 1 \rangle$

- Reverse KL

- $f(u) = -\log u$, $\text{dom}(f) = \mathbb{R}_+$
- $f'(u) = -1/u$
- $f_*(v) = -1 - \log(-v)$, $\text{dom}(f_*) = \mathbb{R}_-$ (c.f. Proposition 2)
- $f'_*(v) = -1/v$
- $D_f(p, q) = \text{KL}(q, p)$

- Jeffreys

- $f(u) = (u - 1) \log u$, $\text{dom}(f) = \mathbb{R}_+$
- $f'(u) = \log u + 1 - \frac{1}{u}$
- $f_*(v) = \frac{1}{W(\exp(1-v))} + \log \frac{1}{W(\exp(1-v))} - 1$, $\text{dom}(f_*) = \mathbb{R}$ (c.f. Proposition 3)
- $f'_*(v) = \frac{1}{W(\exp(1-v))}$
- $D_f(p, q) = \text{KL}(p, q) + \text{KL}(q, p) = \langle \log p - \log q, p - q \rangle$
- Remark: W is the Lambert function

- Jensen-Shannon

- $f(u) = u \log u - (u + 1) \log \left(\frac{u+1}{2} \right)$, $\text{dom}(f) = \mathbb{R}_+$
- $f'(u) = \log u - \log \left(\frac{u+1}{2} \right) = \log \left(\frac{2u}{u+1} \right)$
- $f_*(v) = -\log(2 - \exp(v))$, $\text{dom}(f_*) = (-\infty, \log 2)$ (c.f. Proposition 4)
- $f'_*(v) = \frac{\exp(v)}{2 - \exp(v)} = \frac{1}{2 \exp(-v) - 1}$
- $D_f(p, q) = 2 \cdot \text{JS}(p, q) := \text{KL}(p, m) + \text{KL}(q, m) = \left\langle p, \log \left(\frac{2p}{p+q} \right) \right\rangle + \left\langle q, \log \left(\frac{2q}{p+q} \right) \right\rangle$, where $m := \frac{1}{2}(p + q)$.

- Squared Hellinger

- $f(u) = (\sqrt{u} - 1)^2 = u - 2\sqrt{u} + 1$, $\text{dom}(f) = \mathbb{R}_+$
- $f'(u) = 1 - \frac{1}{\sqrt{u}}$
- $f_*(v) = \frac{v}{1-v}$, $\text{dom}(f_*) = (-\infty, 1)$ (c.f. Proposition 5)
- $f'_*(v) = \frac{1}{(1-v)^2}$
- $D_f(p, q) = 2 \cdot \text{SH}(p, q) = \sum_{j=1}^k (\sqrt{p_j} - \sqrt{q_j})^2$

- Chi-squared divergence (a.k.a. Pearson divergence)

- $f(u) = \frac{1}{2}(u^2 - 1)$, $\text{dom}(f) = \mathbb{R}$

- $f'(u) = u$
- $f_*(v) = \frac{1}{2}(v^2 + 1)$, $\text{dom}(f_*) = \mathbb{R}$ (c.f. Proposition 6)
- $f'_*(v) = v$
- $D_f(p, q) = \frac{1}{2}\chi^2(p, q) := \frac{1}{2} \sum_{j=1}^k \frac{(p_j - q_j)^2}{q_j}$
- Remark: other possible generating functions are $f(u) = \frac{1}{2}(u - 1)^2$ or $f(u) = \frac{1}{2}(u^2 - u)$. We choose $f(u) = \frac{1}{2}(u^2 - 1)$ as it satisfies $f'(0) = 0$.
- Reverse Chi-squared divergence (a.k.a. Neyman divergence)
 - $f(u) = \frac{1}{2}(\frac{1}{u} - 1)$, $\text{dom}(f) = \mathbb{R}_+$
 - $f'(u) = -\frac{1}{2u^2}$
 - $f_*(v) = -\sqrt{-2v} + \frac{1}{2}$, $\text{dom}(f_*) = \mathbb{R}_-$ (c.f. Proposition 7)
 - $f'_*(v) = \frac{1}{\sqrt{-2v}}$
 - $D_f(p, q) = \frac{1}{2}\chi^2(q, p)$
- α -divergences
 - $f(u) = \begin{cases} \frac{(u^\alpha - 1) - \alpha(u - 1)}{\alpha(\alpha - 1)} + \iota_{\mathbb{R}_+}(u) & \text{if } \alpha \neq 1 \\ u \log u - (u - 1) & \text{if } \alpha = 1 \end{cases}$, $\text{dom}(f) = \mathbb{R}_+$
 - $f'(u) = \log_\alpha(u) := \begin{cases} \frac{u^{\alpha-1} - 1}{\alpha - 1} & \text{if } \alpha \neq 1 \\ \log u & \text{if } \alpha = 1 \end{cases}$
 - $f_*(v) = \begin{cases} \frac{(1 + (\alpha - 1)v)^{\frac{\alpha}{\alpha-1}} - 1}{\alpha} & \text{if } \alpha \in (0, 1) \\ \exp(v) & \text{if } \alpha = 1 \\ \frac{(1 + (\alpha - 1)v)_+^{\frac{\alpha}{\alpha-1}} - 1}{\alpha} & \text{if } \alpha \in (1, +\infty), \text{ where } x_+ = \max(x, 0) \end{cases}$, (c.f. Proposition 8)
 - $\text{dom } f_* = \begin{cases} (-\infty, -\frac{1}{\alpha-1}] & \text{if } \alpha \in (0, 1) \\ \mathbb{R} & \text{if } \alpha \geq 1 \end{cases}$
 - $f'_*(v) = \exp_\alpha(v) := \begin{cases} (1 + (\alpha - 1)v)^{\frac{1}{\alpha-1}} & \text{if } \alpha \in (0, 1) \\ \exp v & \text{if } \alpha = 1 \\ (1 + (\alpha - 1)v)_+^{\frac{1}{\alpha-1}} & \text{if } \alpha \in (1, +\infty) \end{cases}$
 - $D_f(p, q) = \frac{1}{\alpha}(\langle p, \log_\alpha(p/q) \rangle - \langle p, 1 \rangle + \langle q, 1 \rangle) = \frac{1}{\alpha(\alpha-1)} (\langle p^\alpha, q^{\alpha-1} \rangle - \alpha \langle p, 1 \rangle + (\alpha-1) \langle q, 1 \rangle)$
 - Remark: By analogy with the KL divergence case, the α -divergence can be derived by setting $f'(u) := \log_\alpha(u)$ and by integrating. The constant of integration is chosen so as to satisfy $f(1) = 0$.

All functions f above can easily be shown to be strictly convex and differentiable. Note that both properties (strict convexity and differentiability) are defined on the interior of the domain of definition of the function, and therefore non-differentiability on its border (typically 0 on the left) does not matter to assess these properties.

C. Proofs

C.1. Conjugates

In this section, we derive the conjugate of f ,

$$f_*(v) := \max_u uv - f(u)$$

for the divergences listed in Appendix B.

Reverse KL divergence.

Proposition 2. *The conjugate of $f(u) = -\log(u)$ for $u \in \text{dom } f = (0, +\infty)$ is $f_*(v) = -1 - \log(-v)$ for $v \in \text{dom } f_* = (-\infty, 0)$.*

Proof. For $v > 0$, the function $u \mapsto uv - f(u)$ is unbounded above by taking $u \rightarrow +\infty$. For $v < 0$, we have $\lim_{u \rightarrow 0} \{uv - f(u)\} = \lim_{u \rightarrow +\infty} \{uv - f(u)\} = -\infty$. Hence a maximizer u^* of $u \mapsto uv - f(u)$ exists and is characterized by $v = f'(u^*)$. Since $f'(u) = -1/u$, we have

$$v = f'(u) \iff v = -1/u \iff u = -1/v = f'_*(v).$$

We then get

$$f_*(v) = -1 + \log(-1/v) = -1 - \log(-v).$$

□

Jeffreys divergence.

Proposition 3. *The conjugate of $f(u) = (u-1)\log u$ for $u \in \text{dom } f = (0, +\infty)$ is $f_*(v) = \frac{1}{W(\exp(1-v))} + \log \frac{1}{W(\exp(1-v))} - 1$ for $v \in \text{dom } f_* = \mathbb{R}$, where W is the Lambert W function.*

Proof. For any $v \in \mathbb{R}$, we have $\lim_{u \rightarrow 0} \{uv - f(u)\} = \lim_{u \rightarrow +\infty} \{uv - f(u)\} = -\infty$. Therefore, for any $v \in \mathbb{R}$, the function $u \mapsto uv - f(u)$ has a maximizer u^* characterized by $v = f'(u^*)$.

Using $f'(u) = \log u + 1 - \frac{1}{u}$, we get

$$v = f'(u) = \log u + 1 - \frac{1}{u} \iff 1 - v = 1/u + \log(1/u) \iff \exp(1-v) = \exp(1/u)1/u.$$

We recall the definition of the Lambert W function.

For $z \geq 0$, $W(z)$ is the inverse of the function $g(w) = w \exp(w)$, i.e., $W(z) = g^{-1}(z)$ and $W^{-1}(w) = g(w)$. We then obtain

$$u = \frac{1}{W(\exp(1-v))} = f'_*(v).$$

Furthermore, we have

$$uv = u \log u + u - 1$$

and therefore, we get

$$\begin{aligned} f_*(v) &= uv - u \log u + \log u \\ &= u + \log u - 1 \\ &= \frac{1}{W(\exp(1-v))} + \log \frac{1}{W(\exp(1-v))} - 1. \end{aligned}$$

□

Jensen-Shannon divergence.

Proposition 4. *The conjugate of $f(u) = u \log u - (u+1) \log \left(\frac{u+1}{2}\right)$ for $u \in \text{dom } f = (0, +\infty)$ is $f_*(v) = \log u - v = -\log(2 - \exp(v))$, for $v \in \text{dom } f_* = (-\infty, \log 2)$.*

Proof. We have for $u \rightarrow +\infty$, $uv - u \log u + (u+1) \log((u+1)/2) \sim uv - u \log u + u \log u/2 \sim uv - u \log 2$, where \sim denotes asymptotic equivalence. Hence for $v > \log 2$, we have $\lim_{u \rightarrow +\infty} uv - f(u) = +\infty$. For $v < \log 2$, we have $\lim_{u \rightarrow +\infty} \{uv - f(u)\} = \lim_{u \rightarrow 0} \{uv - f(u)\} = -\infty$. Hence, for $v < \log 2$, a maximizer u^* of $u \mapsto uv - f(u)$ exists and is characterized by $v = f'(u^*)$.

Using $f'(u) = \log\left(\frac{2u}{u+1}\right)$, we get

$$v = f'(u) \iff \exp(v) = \frac{2u}{u+1} \iff u = \frac{\exp(v)}{2 - \exp(v)} = \frac{\exp(0)}{2\exp(-v) - \exp(0)} = f'_*(v).$$

Let us now compute $f_*(v) = uv - u \log u + (u+1) \log\left(\frac{u+1}{2}\right)$. We observe that

$$u+1 = \frac{2}{2 - \exp(v)} \iff \frac{u+1}{2} = \frac{u}{\exp(v)} \iff \log\left(\frac{u+1}{2}\right) = \log u - v.$$

Therefore

$$(u+1) \log\left(\frac{u+1}{2}\right) = u(\log u - v) + \log u - v = u \log u - uv + \log u - v.$$

Therefore

$$f_*(v) = \log u - v = -\log(2 - \exp(v)).$$

□

Squared Hellinger divergence.

Proposition 5. *The conjugate of $f(u) = (\sqrt{u} - 1)^2$ for $u \in \text{dom } f = [0, +\infty)$ is $f_*(v) = \frac{v}{1-v}$ for $v \in \text{dom } f_* = (-\infty, 1)$.*

Proof. We have $uv - f(u) = u(v-1) + 2\sqrt{u} - 1$. Hence, for $v \geq 1$, $u \rightarrow uv - f(u)$ is unbounded above. For $v < 1$, $\lim_{u \rightarrow +\infty} \{uv - f(u)\} = -\infty$ and $uv - f(u)$ has then a maximizer u^* in $[0, +\infty)$ characterized by $f'(u^*) = v$. We have $f(u) = (\sqrt{u} - 1)^2 = u - 2\sqrt{u} + 1$. Using

$$v = f'(u) = 1 - \frac{1}{\sqrt{u}} \iff \sqrt{u} = \frac{1}{1-v} \iff u = \frac{1}{(1-v)^2},$$

we get

$$\begin{aligned} uv &= \frac{v}{(1-v)^2} \\ \sqrt{u} - 1 &= \frac{v}{1-v} \\ (\sqrt{u} - 1)^2 &= \frac{v^2}{(1-v)^2} \end{aligned}$$

so that

$$f^*(v) = \frac{v}{(1-v)^2} - \frac{v^2}{(1-v)^2} = \frac{v}{1-v}.$$

□

Chi-squared divergence (Pearson divergence).

Proposition 6. *The conjugate of $f(u) = \frac{1}{2}(u^2 - 1)$ for $u \in \text{dom } f = \mathbb{R}$ is $f_*(v) = \frac{1}{2}(v^2 + 1)$ for $v \in \text{dom } f_* = \mathbb{R}$.*

Proof. f is a strongly convex function on \mathbb{R} . Hence the maximizer u^* of $u \mapsto uv - f(u)$ exists for any $v \in \mathbb{R}$ and is characterized by $f'(u^*) = v$. Using $f'(u) = u$, we get $v = f'(u) = u$. Therefore,

$$f_*(v) = uv - f(u) = v^2 - \frac{1}{2}(v^2 - 1) = \frac{1}{2}(v^2 + 1).$$

□

Reverse Chi-squared divergence (Neyman divergence).

Proposition 7. The conjugate of $f(u) = \frac{1}{2}(\frac{1}{u} - 1)$ for $u \in \text{dom } f = (0, +\infty)$ is $f_*(v) = -\sqrt{-2v} + \frac{1}{2}$ for $v \in \text{dom } f_* = (-\infty, 0)$.

Proof. For $v > 0$, $\lim_{u \rightarrow +\infty} \{uv - f(u)\} = +\infty$. For $v < 0$, $\lim_{u \rightarrow 0} \{uv - f(u)\} = \lim_{u \rightarrow +\infty} \{uv - f(u)\} = -\infty$. Hence the maximizer u^* of $u \mapsto uv - f(u)$ exists and is characterized by $f'(u^*) = v$. Using $v = f'(u) = -\frac{1}{2u^2}$, we get

$$u^2 = -\frac{1}{2v} \iff u = \frac{1}{\sqrt{-2v}}.$$

Therefore,

$$f_*(v) = uv - \frac{1}{2} \left(\frac{1}{u} - 1 \right) = \frac{v}{\sqrt{-2v}} - \frac{\sqrt{-2v}}{2} + \frac{1}{2} = -\sqrt{-2v} + \frac{1}{2}$$

□

α -divergences.

Proposition 8. The conjugate of $f(u) = \frac{(u^\alpha - 1) - \alpha(u - 1)}{\alpha(\alpha - 1)} + \iota_{\mathbb{R}_+}(u)$ is

$$f_*(v) = \begin{cases} \frac{(1 + (\alpha - 1)v)^{\frac{\alpha}{\alpha - 1}} - 1}{\alpha} & \text{if } \alpha \in (0, 1) \\ \exp(v) & \text{if } \alpha = 1 \\ \frac{(1 + (\alpha - 1)v)^{\frac{\alpha}{\alpha - 1}} - 1}{\alpha} & \text{if } \alpha \in (1, +\infty), \text{ where } x_+ = \max(x, 0) \end{cases}$$

for

$$v \in \text{dom } f_* = \begin{cases} (-\infty, -\frac{1}{\alpha - 1}] & \text{if } \alpha \in (0, 1) \\ \mathbb{R} & \text{if } \alpha \geq 1 \end{cases}$$

Proof. **Case** $\alpha \in (0, 1)$. As $u \rightarrow +\infty$, we have $vu - \frac{(u^\alpha - 1) - \alpha(u - 1)}{\alpha(\alpha - 1)} \sim \left(v + \frac{1}{\alpha - 1}\right)u$.

Hence for $v > \frac{1}{1 - \alpha}$, $\lim_{u \rightarrow +\infty} vu - \frac{(u^\alpha - 1) - \alpha(u - 1)}{\alpha(\alpha - 1)} = +\infty$.

For $v \leq \frac{1}{1 - \alpha}$, $\lim_{u \rightarrow +\infty} \{vu - \frac{(u^\alpha - 1) - \alpha(u - 1)}{\alpha(\alpha - 1)}\} = \lim_{u \rightarrow 0} \{vu - \frac{(u^\alpha - 1) - \alpha(u - 1)}{\alpha(\alpha - 1)}\} = -\infty$. Hence the maximizer u^* of $u \mapsto uv - f(u)$ exists and is characterized by $f'(u^*) = v$.

We have

$$v = f'(u) = \frac{u^{\alpha - 1} - 1}{\alpha - 1} \iff u = (1 + (\alpha - 1)v)^{\frac{1}{\alpha - 1}} = f'_*(v)$$

Therefore, denoting $z = (1 + (\alpha - 1)v)$, we have, for $v \leq \frac{1}{1 - \alpha}$

$$\begin{aligned} f^*(v) &= uv - f(u) \\ &= vz^{\frac{1}{\alpha - 1}} - \frac{(z^{\frac{\alpha}{\alpha - 1}} - 1) - \alpha(z^{\frac{1}{\alpha - 1}} - 1)}{\alpha(\alpha - 1)} \\ &= \frac{\alpha((\alpha - 1)v + 1)z^{\frac{1}{\alpha - 1}} - z^{\frac{\alpha}{\alpha - 1}} - (\alpha - 1)}{\alpha(\alpha - 1)} \\ &= \frac{z^{\frac{\alpha}{\alpha - 1}} - 1}{\alpha} \\ &= \frac{(1 + (\alpha - 1)v)^{\frac{\alpha}{\alpha - 1}} - 1}{\alpha}. \end{aligned}$$

Case $\alpha \in (1, +\infty)$. For any $v \in \mathbb{R}$, $\lim_{u \rightarrow +\infty} vu - \frac{(u^\alpha - 1) - \alpha(u - 1)}{\alpha(\alpha - 1)} = -\infty$. Denote $g(u) = vu - \frac{(u^\alpha - 1) - \alpha(u - 1)}{\alpha(\alpha - 1)}$. Note that g' is decreasing on $[0, +\infty]$.

If $v \leq \frac{1}{1-\alpha}$, we have $g'(0) = v + \frac{1}{\alpha-1} \leq 0$, hence $g'(u) < 0$ for any $u > 0$, that is g is decreasing on $[0, +\infty)$ and $\sup_{u \geq 0} g(u) = g(0) = -f(0)$ in this case.

If $v > \frac{1}{1-\alpha}$, $g'(0) > 0$, $\lim_{u \rightarrow +\infty} g'(u) = -\infty$ and a maximizer u^* of g can be found by the necessary condition $f'(u^*) = v$. Following the computations done in the case $\alpha \in (0, 1)$, for $v > \frac{1}{1-\alpha}$ we get

$$\begin{aligned} f^*(v) &= \frac{(1 + (\alpha - 1)v)^{\frac{\alpha}{\alpha-1}}}{\alpha} \\ f'_*(v) &= (1 + (\alpha - 1)v)^{\frac{1}{\alpha-1}}. \end{aligned}$$

And for any $v \in \mathbb{R}$, we get

$$\begin{aligned} f^*(v) &= \frac{(1 + (\alpha - 1)v)_+^{\frac{\alpha}{\alpha-1}} - 1}{\alpha} \\ f'_*(v) &= (1 + (\alpha - 1)v)_+^{\frac{1}{\alpha-1}}. \end{aligned}$$

where $x_+ = \max\{x, 0\}$.

Case $\alpha = 1$ Follows similar computations as for the KL case. □

C.2. f -softargmaxes

KL divergence. With $f_*(v) = f'_*(v) = \exp(v - 1)$, we obtain

$$\tau^* + 1 = \log \sum_{j=1}^k q_j \exp(\theta_j).$$

Plugging τ^* back in (13), we obtain

$$p_j^* = \frac{q_j \exp(\theta_j)}{\sum_{j=1}^k q_j \exp(\theta_j)},$$

Other choices of f do not lead to closed-form expressions.

C.3. f -softplus and f -sigmoid

In this section, we derive closed-form expressions for some instances of f -softplus and f -sigmoid. In this binary classification setting, we define $\theta := (0, s)$, $\mathbf{q} = (q_0, q_1)$ and $\mathbf{p}^* = (1 - \pi^*, \pi^*)$. By Proposition 1, we seek for

$$\tau^* = \operatorname{argmin}_{\tau \in \mathbb{R}} \tau + q_0 f'_*(-\tau) + q_1 f'_*(s - \tau).$$

Since f_* differentiable and convex, such a scalar τ^* is characterized by

$$-\tau^* \in \operatorname{dom} f_*, \tag{16}$$

$$s - \tau^* \in \operatorname{dom} f_*, \tag{17}$$

$$q_0 f'_*(-\tau^*) + q_1 f'_*(s - \tau^*) = 1. \tag{18}$$

KL divergence.

Proposition 9. Let $f(u) = u \log u$. Then,

$$\operatorname{sigmoid}_f(s; \mathbf{q}) = \frac{q_1 \exp(s)}{q_0 + q_1 \exp(s)}.$$

Proof. Using (18) and $f'_*(v) = \exp(v - 1)$, we obtain

$$\tau^* + 1 = \log(q_0 + q_1 \exp(s)).$$

so that

$$\pi^* = \frac{q_1 \exp(s)}{q_0 + q_1 \exp(s)}.$$

□

Reverse KL divergence.

Proposition 10. *Let $f(u) = -\log(u)$. Then,*

$$\begin{aligned} \text{softplus}_f(s; \mathbf{q}) &= \tau^* - q_0 \log(\tau^*) - q_1 \log(\tau^* - s) - (q_0 + q_1) \\ \text{sigmoid}_f(s; \mathbf{q}) &= \frac{q_1}{\tau^* - s} \end{aligned}$$

where

$$\tau^* := \frac{1}{2} \left(q_0 + q_1 + s + \sqrt{(q_0 + q_1 + s)^2 - 4q_0 s} \right).$$

Proof. Using $f'_*(v) = -1/v$, the root condition (18) can be written as

$$\begin{aligned} & q_0 f'_*(-\tau) + q_1 f'_*(s - \tau) = 1 \\ \iff & \frac{q_0}{\tau} + \frac{q_1}{\tau - s} = 1 \\ \iff & \frac{q_0(\tau - s) + q_1 \tau}{\tau(\tau - s)} = 1 \\ \iff & a\tau^2 + b\tau + c = 0, \end{aligned}$$

for

$$\begin{aligned} a &:= 1 \\ b &:= -(q_0 + q_1 + s) \\ c &:= q_0 s. \end{aligned}$$

Let us define the discriminant

$$\begin{aligned} \Delta &:= b^2 - 4ac \\ &= (q_0 + q_1 + s)^2 - 4q_0 s \\ &= (q_1 - q_0 + s)^2 + 4q_0 q_1 \\ &> 0. \end{aligned}$$

Therefore, the root condition (18) reads

$$\begin{aligned} \tau^* &\in \{\tau_1, \tau_2\} \\ &:= \left\{ \frac{-b + \sqrt{\Delta}}{2a}, \frac{-b - \sqrt{\Delta}}{2a} \right\}. \end{aligned}$$

Since $\text{dom } f_* = \mathbb{R}_-$, we have

$$(-\tau^* \in \text{dom } f_*) \text{ and } (s - \tau^* \in \text{dom } f_*) \iff \tau^* \geq \max\{0, s\}.$$

We have

$$\begin{aligned} \tau_2 &= \frac{1}{2} \left(q_0 + q_1 + s - \sqrt{(q_0 + q_1 + s)^2 - 4q_0 s} \right) \\ \tau_2 - s &= \frac{1}{2} \left(q_0 + q_1 - s - \sqrt{(q_0 + q_1 - s)^2 + 4q_1 s} \right). \end{aligned}$$

Hence for $s < 0$, $\tau_2 < 0$, and for $s \geq 0$, $\tau_2 - s < 0$. So the unique solution to the set of conditions (16), (17) and (18) is

$$\tau^* = \tau_1 = \frac{1}{2} \left(q_0 + q_1 + s + \sqrt{(q_0 + q_1 + s)^2 - 4q_0s} \right)$$

Plugging this value in the formulas for the softplus and the sigmoid in Proposition 1 gives the result. \square

Numerically stable implementation.

For very large positive or negative s we may get numerical issues of the form $as - bs$. For a numerically stable implementation, one could write

$$\begin{aligned} \tau^* &= q_0 + h(g(s)) \\ \tau^* - s &= q_1 + h(-g(s)) \end{aligned}$$

where

$$\begin{aligned} g(s) &:= q_1 - q_0 + s \\ h(g) &:= \frac{1}{2} \cdot \begin{cases} |g| \left(\sqrt{1 + \frac{4q_0q_1}{g^2}} - 1 \right) & \text{if } g < -1 \\ g + \sqrt{g^2 + 4q_0q_1} & \text{otherwise} \end{cases} \end{aligned}$$

We can then further rewrite

$$\begin{aligned} |g| \left(\sqrt{1 + \frac{4q_0q_1}{g^2}} - 1 \right) &= |g| \frac{4q_0q_1/g^2}{\sqrt{1 + \frac{4q_0q_1}{g^2}} + 1} \\ &= \frac{4q_0q_1}{\sqrt{g^2 + 4q_0q_1} + |g|}. \end{aligned}$$

Jensen-Shannon divergence.

Proposition 11. Let $f(u) = u \log u - (u + 1) \log \left(\frac{u+1}{2} \right)$. Then,

$$\begin{aligned} \text{softplus}_f(s; \mathbf{q}) &= \log(x^*) - q_0 \log(2 - 1/x^*) - q_1 \log(2 - 1/(x^* \cdot y)) \\ \text{sigmoid}_f(s; \mathbf{q}) &= \frac{q_1}{2 \cdot x^* \cdot y - 1}, \end{aligned}$$

where

$$\begin{aligned} x^* &:= \frac{-b + \sqrt{\Delta}}{2a} \\ y &:= \exp(-s) \\ a &:= 4y \\ b &:= -2(1 + y + yq_0 + q_1) \\ c &:= 1 + q_0 + q_1 \\ \Delta &:= b^2 - 4ac. \end{aligned}$$

Proof. Using $f'_*(v) = \frac{1}{2 \exp(-v) - 1}$, the root equation (18)

$$\frac{q_0}{2 \exp(\tau) - 1} + \frac{q_1}{2 \exp(\tau - s) - 1} = 1.$$

Using the change of variables $x := \exp(\tau)$ and $y := \exp(-s)$, we obtain

$$\begin{aligned} \frac{q_0}{2x - 1} + \frac{q_1}{2xy - 1} &= 1 \\ \iff q_0(2xy - 1) + q_1(2x - 1) &= (2x - 1)(2xy - 1) \\ \iff ax^2 + bx + c &= 0, \end{aligned}$$

where we defined

$$\begin{aligned} a &:= 4y \\ b &:= -2(1 + y + yq_0 + q_1) \\ c &:= 1 + q_0 + q_1. \end{aligned}$$

Let us define the discriminant

$$\Delta := b^2 - 4ac.$$

We have

$$\begin{aligned} \Delta &= 4y \left[\left((1 + q_0)y^{1/2} + (1 + q_1)y^{-1/2} \right)^2 - 4(1 + q_0 + q_1) \right] \\ &\geq 4y [4(1 + q_0)(1 + q_1) - 4(1 + q_0 + q_1)] \\ &= 16yq_0q_1 \\ &> 0, \end{aligned}$$

where in the second line we used that $y > 0$ and $\min_{x>0} \alpha x + \beta x^{-1} = 2\sqrt{\alpha\beta}$ for any $\alpha > 0, \beta > 0$. Therefore, we have

$$x^* = \exp(\tau^*) = \frac{-b + \sqrt{\Delta}}{2a}.$$

Using (12) with

$$\begin{aligned} f_*(-\tau^*) &= -\log(2 - \exp(-\tau^*)) = -\log(2 - 1/x^*) \\ f_*(s - \tau^*) &= -\log(2 - \exp(s - \tau^*)) = -\log(2 - 1/(x^* \cdot y)), \end{aligned}$$

we obtain

$$\text{softplus}_f(s) = \log(x^*) - q_0 \log(2 - 1/x^*) - q_1 \log(2 - 1/(x^* \cdot y)).$$

Similarly, using (13), we obtain

$$\text{sigmoid}_f(s) = \frac{q_1}{2 \exp(\tau^* - s) - 1} = \frac{q_1}{2 \cdot x^* \cdot y - 1}.$$

□

Numerically stable implementation.

Working with exponential requires careful handling of any formula like $uy - vy$ or $(uy)/(vy)$ that can easily reduce numerically to $\infty - \infty = \text{NaN}$ or $\infty/\infty = \text{NaN}$. We therefore derive the expressions of x^* and x^*y in terms of numerically stable operations.

We have, denoting $\alpha := 1 + q_0$ and $\beta := 1 + q_1$,

$$\begin{aligned} x^* &= \frac{\alpha + \beta y^{-1}}{4} + \frac{1}{4} \sqrt{(\alpha + \beta y^{-1})^2 - 4cy^{-1}} \\ &= \frac{y^{-1/2}}{4} \left(\alpha y^{1/2} + \beta y^{-1/2} + \sqrt{(\alpha y^{1/2} + \beta y^{-1/2})^2 - 4c} \right) \\ &= \frac{y^{-1/2}}{4} (\alpha y^{1/2} + \beta y^{-1/2}) \left(1 + \sqrt{1 - \frac{4(1 + q_0 + q_1)}{(\alpha y^{1/2} + \beta y^{-1/2})^2}} \right) \\ &= \exp(s/2 + h(s/2)), \end{aligned}$$

where

$$\begin{aligned} h(s) &:= g(s) + \log \left(1 + \sqrt{1 - \frac{4(1 + q_0 + q_1)}{\exp(2g(s))}} \right) - 2 \log(2) \\ g(s) &:= \log(\alpha e^{-s} + \beta e^s) \\ &= \text{logsumexp}(-s + \log(1 + q_0), s + \log(1 + q_1)). \end{aligned}$$

To summarize, we have

$$\begin{aligned} x^\star &= \exp(s/2 + h(s/2)) \\ x^\star y &= \exp(-s/2 + h(s/2)) \end{aligned}$$

and therefore

$$\begin{aligned} \text{softplus}_f(s) &= s/2 + h(s/2) - q_0 \log[2 - \exp(-s/2 - h(s/2))] - q_1 \log[2 - \exp(s/2 - h(s/2))] \\ \text{sigmoid}_f(s) &= \frac{q_1}{2 \exp(-s/2 + h(s/2)) - 1}. \end{aligned}$$

Squared Hellinger divergence. Using $f'_*(v) = \frac{1}{(1-v)^2}$ and using the change of variable $x := \tau + 1$, the root equation (18) becomes

$$\begin{aligned} \frac{q_0}{(1+\tau)^2} + \frac{q_1}{(1+\tau-s)^2} = 1 &\iff \frac{q_0}{x^2} + \frac{q_1}{(x-s)^2} = 1 \\ &\iff q_0(x-s)^2 + q_1x^2 = x^2(x-s)^2 \\ &\iff q_0(x^2 - 2sx + s^2) + q_1x^2 = x^2(x^2 - 2sx + s^2) \\ &\iff q_0(x^2 - 2sx + s^2) + q_1x^2 = x^4 - 2sx^3 + s^2x^2 \\ &\iff ax^4 + bx^3 + cx^2 + dx + e = 0, \end{aligned}$$

where

$$\begin{aligned} a &:= 1 \\ b &:= -2s \\ c &:= s^2 - q_0 - q_1 \\ d &:= 2sq_0 \\ e &:= -q_0s^2. \end{aligned}$$

This is a quartic equation, which can be solved in closed form.

C.4. Proof of Proposition 1

First, the maximum defining the softmax and softargmax is well defined since it is a strictly concave problem on a non-empty bounded set. Denoting $\Omega_j(p) := q_j f(p/q_j)$, such that $\Omega'_j(p) = f'(p/q_j)$ and $\Omega_j^*(\theta) = q_j f^*(\theta)$, we can apply Lemma 1 and get

$$\max_{\mathbf{p} \in \Delta^k} \langle \mathbf{p}, \boldsymbol{\theta} \rangle - \sum_{j=1}^k q_j f(p_j/q_j) = \inf_{\tau \in \mathbb{R}} \tau + \sum_{j=1}^k q_j f^*(\max\{\theta_j - \tau, f'(0)\}), \quad (19)$$

with $f'(0) := \lim_{x \rightarrow 0, x \geq 0} f'(x) \in \mathbb{R} \cup \{-\infty\}$.

Since $(0, +\infty) \subseteq \text{dom } f'$, and $\mathbf{q} > 0$, $f' \left(\left(\sum_{j=1}^k q_j \right)^{-1} \right)$ and $f'(q_{j_{\max}}^{-1})$ for $j_{\max} \in \arg\max_{j \in \{1, \dots, k\}} \theta_j$ are well defined.

We can then define

$$\begin{aligned} \tau_{\min} &:= \theta_{\max} - f'(q_{j_{\max}}^{-1}) \\ \tau_{\max} &:= \theta_{\max} - f' \left(\left(\sum_{j=1}^k q_j \right)^{-1} \right) \end{aligned}$$

where $\theta_{\max} = \theta_{j_{\max}}$. Since $\mathbf{q} > 0$, $q_{j_{\max}} < \sum_{j=1}^k q_j$, and since f' is increasing, we have $\tau_{\min} < \tau_{\max}$.

We can then analyze the following function on $[\tau_{\min}, \tau_{\max}]$:

$$h(\tau) := \tau + \sum_{j=1}^k q_j f^*(\max\{\theta_j - \tau, f'(0)\}).$$

First, we need to ensure that we can compute derivatives of this function in $[\tau_{\min}, \tau_{\max}]$. From Lemma 2, we have $f^*(\max\{\theta_j - \tau, f'(0)\}) = (f + \iota_{\mathbb{R}_+})^*(\theta_j - \tau)$ and $\text{dom}(f + \iota_{\mathbb{R}_+})'_* = \text{dom } f'_* \cup (-\infty, f'(0)]$. Therefore, if $f'(0) > -\infty$, the domain of $(f + \iota_{\mathbb{R}_+})'_*$ is unbounded below. Otherwise, if $f'(0) = -\infty$, since $\text{im } f' \subseteq \text{dom } f'_*$, the domain of f'_* , and so of $(f + \iota_{\mathbb{R}_+})'_*$, are unbounded below. Denoting $\alpha := \sup \text{dom}(f + \iota_{\mathbb{R}_+})'_*$, since $\theta_{\max} - \tau_{\min} = f'(q_{j_{\max}}^{-1}) \in \text{dom } f'_*$, we then have $\theta_{\max} - \tau_{\min} < \alpha$ and therefore

$$\begin{aligned} \tau &\geq \tau_{\min} \\ \implies \theta_{\max} - \tau &< \alpha \\ \iff \theta_j - \tau &< \alpha, \text{ for all } j \in \{1, \dots, k\} \\ \implies \tau &\in \text{dom } h'. \end{aligned}$$

We can then show that $h'(\tau_{\min}) \leq 0$ and $h'(\tau_{\max}) \geq 0$. Indeed, we have

$$\begin{aligned} h'(\tau_{\min}) &= 1 - \sum_{j=1}^k q_j f'_*(\max\{\theta_j - \tau_{\min}, f'(0)\}) \\ &\stackrel{(i)}{\leq} 1 - q_{j_{\max}} f'_*(\max\{\theta_{\max} - \tau_{\min}, f'(0)\}) \\ &= 1 - q_{j_{\max}} f'_*(\max\{f'(q_{j_{\max}}^{-1}), f'(0)\}) \\ &\stackrel{(ii)}{=} 0, \end{aligned}$$

where in (i) we used that $q > 0$ and $f'_*(\max\{y, f'(0)\}) \geq 0$ for any $y \in \text{dom}(f + \iota_{\mathbb{R}_+})'_*$ as per Lemma 2, and in (ii), we used that $q_{j_{\max}}^{-1} > 0$, f' is increasing and $f'_*(f'(p)) = p$ for any $p \in \text{dom } f'$.

Similarly, we have

$$\begin{aligned} h'(\tau_{\max}) &= 1 - \sum_{j=1}^k q_j f'_*(\max\{\theta_j - \tau_{\max}, f'(0)\}) \\ &\stackrel{(i)}{\geq} 1 - \left(\sum_{j=1}^k q_j\right) f'_*(\max\{\theta_{\max} - \tau_{\max}, f'(0)\}) \\ &= 1 - \left(\sum_{j=1}^k q_j\right) f'_*\left(\max\left\{f'\left(\left(\sum_{j=1}^k q_j\right)^{-1}\right), f'(0)\right\}\right) \\ &\stackrel{(ii)}{=} 0, \end{aligned}$$

where in (i) we used that $\sum_j a_j b_j \leq (\sum_j a_j) \max_j b_j$ if $a_j \geq 0$ with here $a_j = q_j > 0$ and $b_j = f'_*(\max\{\theta_j - \tau_{\max}, f'(0)\})$, and in (ii) we used the same reasoning as for $h'(\tau_{\min})$.

Finally, we show that h' is increasing on $[\tau_{\min}, \tau_{\max}]$. For $j = j_{\max}$, we have that

$$\begin{aligned} \tau &\leq \tau_{\max} \\ \iff \theta_{\max} - \tau &\geq \theta_{\max} - \tau_{\max} = f'(q_{j_{\max}}) > f'(0). \end{aligned}$$

Since $f'_*(\max\{\theta, f'(0)\}) = (f + \iota_{\mathbb{R}_+})^*$ is increasing on $\text{dom } f^* \setminus (-\infty, f(0)]$ (see Proposition 2), $\tau \mapsto -f'_*(\max\{\theta_{\max} - \tau, f'(0)\})$ is increasing on $[\tau_{\min}, \tau_{\max}]$ and so $h'(\tau) = 1 - \sum_{j=1}^k q_j f'_*(\max\{\theta_j - \tau, f'(0)\})$ is increasing on $[\tau_{\min}, \tau_{\max}]$ (as a sum of an increasing and non decreasing function).

Overall, h is well defined and strictly convex on $[\tau_{\min}, \tau_{\max}]$ such that $h'(\tau_{\min}) \leq 0$ and $h'(\tau_{\max}) \geq 0$. Hence, we have

$$\inf_{\tau \in \mathbb{R}} h(\tau) = \min_{\tau_{\min} \leq \tau \leq \tau_{\max}} h(\tau).$$

and the unique minimizer τ^* can then be found by solving the first order optimality condition $h'(\tau^*) = 0$ in $[\tau_{\min}, \tau_{\max}]$. This gives the expression of τ^* in the claim and plugging τ^* back into (19) gives the expression of the softmax. The expression of the softargmax follows from Lemma 1.

C.5. Lemmas

Lemma 1. Given k strictly convex differentiable univariate scalar functions Ω_j such that $(0, +\infty) \subseteq \text{dom } \Omega_j$, for any $\theta \in \mathbb{R}^k$, we have

$$\sup_{\mathbf{p} \in \Delta^k} \langle \mathbf{p}, \theta \rangle - \sum_{j=1}^k \Omega_j(p_j) = \inf_{\tau \in \mathbb{R}} \tau + \sum_{j=1}^k (\Omega_j + \iota_{\mathbb{R}_+})^*(\theta_j - \tau),$$

where

$$(\Omega_j + \iota_{\mathbb{R}_+})^*(z) = \Omega_j^*(\max\{\theta_j - \tau, \Omega_j'(0)\})$$

and $\Omega_j'(0) := \lim_{x \rightarrow 0, x \geq 0} \Omega_j'(x) \in \mathbb{R} \cup \{-\infty\}$. Given a minimizer τ^* of the right hand-side, the maximizer \mathbf{p}^* of the left hand side is given by

$$p_j^* = (\Omega_j + \iota_{\mathbb{R}_+})'_*(\theta_j - \tau^*) = (\Omega_j^*)'(\max\{\theta_j - \tau^*, \Omega_j'(0)\}).$$

Proof. We have

$$\begin{aligned} \sup_{\mathbf{p} \in \Delta^k} \langle \mathbf{p}, \theta \rangle - \sum_{j=1}^k \Omega_j(p_j) &= \sup_{\mathbf{p} \in \mathbb{R}_+^k} \inf_{\tau \in \mathbb{R}} \langle \mathbf{p}, \theta \rangle - \sum_{j=1}^k \Omega_j(p_j) + (1 - \langle \mathbf{p}, \mathbf{1} \rangle) \tau \\ &= \inf_{\tau \in \mathbb{R}} \tau + \max_{\mathbf{p} \in \mathbb{R}_+^k} \langle \mathbf{p}, \theta - \tau \mathbf{1} \rangle - \sum_{j=1}^k \Omega_j(p_j) \\ &= \inf_{\tau \in \mathbb{R}} \tau + \sum_{j=1}^k (\Omega_j + \iota_{\mathbb{R}_+})^*(\theta_j - \tau), \end{aligned}$$

where the second equality stands from strong duality, using that the simplex is convex, that the Ω_j are convex functions, and that the maximization problem is strictly feasible since $(0, +\infty) \subseteq \text{dom } \Omega_j$.

We can then apply Lemma 2, and get

$$\max_{\mathbf{p} \in \Delta^k} \langle \mathbf{p}, \theta \rangle - \sum_{j=1}^k \Omega_j(p_j) = \inf_{\tau \in \mathbb{R}} \tau + \sum_{j=1}^k \Omega_j^*(\max\{\theta_j - \tau, \Omega_j'(0)\}),$$

with $\Omega_j'(0) = \lim_{x \rightarrow 0, x \geq 0} \Omega_j'(x) \in \mathbb{R} \cup \{-\infty\}$. Given a minimizer τ^* of the right hand-side, the maximizer \mathbf{p}^* of the left hand side is given by strong duality as

$$p_j^* = \operatorname{argmax}_{p_j \geq 0} p_j(\theta_j - \tau) - \Omega_j(p_j),$$

whose full expression follows from Lemma 2. □

Lemma 2. Given a strictly convex differentiable univariate scalar function f such that $(0, +\infty) \subseteq \text{dom } f$, we have, for any $y \in \mathbb{R}$,

$$(f + \iota_{\mathbb{R}_+})^*(y) := \sup_{x \geq 0} xy - f(x) = f^*(\max\{y, f'(0)\})$$

with $f'(0) := \lim_{x \rightarrow 0, x \geq 0} f'(x) \in \mathbb{R} \cup \{-\infty\}$. For $y \in \text{dom}(f + \iota_{\mathbb{R}_+})'_* = \text{dom } f'_* \cup (-\infty, f'(0)]$, we have

$$(f + \iota_{\mathbb{R}_+})'_*(y) = f'_*(\max\{y, f'(0)\}) \geq 0.$$

Finally, $(f + \iota_{\mathbb{R}_+})'_*$ is increasing on $(\text{dom } f + \iota_{\mathbb{R}_+})'_* \setminus (-\infty, f'(0)]$.

Proof. Fix some $y \in \mathbb{R}$, denote $f'(0) = \lim_{x \rightarrow 0, x \geq 0} f'(x) \in \mathbb{R} \cup \{-\infty\}$ and

$$h(x) := xy - f(x).$$

If $y \leq f'(0)$ (so provided that $f'(0) > -\infty$), then $h'(x) = y - f'(x) > 0$ for all $x > 0$, since f' is increasing. So h is decreasing on $(0, +\infty)$ and the maximum of h on \mathbb{R}_+ is reached at 0, that is

$$\sup_{x \geq 0} xy - f(x) = -f(0).$$

If $y \geq f'(0)$, then $h'(x) = y - f'(x) > 0$ for all $x < 0$ since f' is increasing. Therefore $h(x) < h(0)$ for all $x < 0$. Finally, since $h'(0) \leq 0$, the supremum of h on \mathbb{R} is necessarily greater or equal than 0, that is,

$$\sup_{x \geq 0} xy - f(x) = \sup_{x \in \mathbb{R}} xy - f(x) = f^*(y).$$

Note that for $y = f'(0)$ (provided that $f'(0) > -\infty$), we then get

$$\sup_{x \geq 0} xf'(0) - f(x) = -f(0) = f^*(f'(0)),$$

and

$$f'_*(f'(0)) = \operatorname{argmax}_{x \in \mathbb{R}} \{xf'(0) - f(x)\} = 0.$$

Combining the two cases above, we get

$$\begin{aligned} \sup_{x \geq 0} xy - f(x) &= \begin{cases} f^*(y) & \text{if } y > f'(0) \\ -f(0) & \text{if } y \leq f'(0) \end{cases} \\ &= f^*(\max\{y, f'(0)\}). \end{aligned}$$

For the derivative, first note that since $f + \iota_{\mathbb{R}_+}$ is strictly convex, its convex conjugate is differentiable on its domain of definition. Then, given $y \in \operatorname{dom}(f + \iota_{\mathbb{R}_+})'_*$, we have from the previous considerations,

$$(f + \iota_{\mathbb{R}_+})'_*(y) = \operatorname{argmax}_{x \geq 0} \{xy - f(x)\} = \begin{cases} f'_*(y) & \text{if } y > f'(0) \\ 0 & \text{if } y \leq f'(0) \end{cases}.$$

Using that $f'_*(f'(0)) = 0$, we get

$$(f + \iota_{\mathbb{R}_+})'_*(y) = f'_*(\max\{y, f'(0)\}).$$

The expressions above also show that $\operatorname{dom}(f + \iota_{\mathbb{R}_+})'_* = \operatorname{dom} f'_* \cup (-\infty, f'(0)]$.

In any case, for $y \in \operatorname{dom}((f + \iota_{\mathbb{R}_+})^*)'$,

$$(f + \iota_{\mathbb{R}_+})'_*(y) = \operatorname{argmax}_{x \geq 0} xy - f(x) \geq 0,$$

by definition of the maximization set.

Finally, since f is strictly convex and differentiable on $(0, +\infty)$, f' is invertible on $(f'(0), f'(\infty))$ where $f'(0) = \lim_{x \rightarrow 0, x \geq 0} f'(x)$ and $f'(\infty) = \lim_{x \rightarrow +\infty} f'(x)$. Moreover, for any $y \in (f'(0), f'(\infty))$, we then have $(f + \iota_{\mathbb{R}_+})'_*(y) = f'_*(y) = (f')^{-1}(y)$, where $(f')^{-1}$ is the inverse of f' on $(f'(0), f'(\infty))$. Since f' is increasing on $(0, +\infty)$, $(f')^{-1}$ is also increasing on $(f'(0), f'(\infty))$, and so are f'_* and $(f + \iota_{\mathbb{R}_+})'_*$. Noting that $(f'(0), f'(\infty)) = \operatorname{dom} f'_* \setminus (-\infty, f'(0)] = (\operatorname{dom} f + \iota_{\mathbb{R}_+})'_* \setminus (-\infty, f'(0)]$ concludes the claim. \square