
Learning Energy Networks with Generalized Fenchel-Young Losses

Mathieu Blondel, Felipe Llinares-López,
Robert Dadashi, Léonard Hussenot, Matthieu Geist
Google Research, Brain team
{mblondel, fllinares, dadashi, hussenot, mfgeist}@google.com

Abstract

Energy-based models, a.k.a. energy networks, perform inference by optimizing an energy function, typically parametrized by a neural network. This allows one to capture potentially complex relationships between inputs and outputs. To learn the parameters of the energy function, the solution to that optimization problem is typically fed into a loss function. The key challenge for training energy networks lies in computing loss gradients, as this typically requires argmin/argmax differentiation. In this paper, building upon a generalized notion of conjugate function, which replaces the usual bilinear pairing with a general energy function, we propose generalized Fenchel-Young losses, a natural loss construction for learning energy networks. Our losses enjoy many desirable properties and their gradients can be computed efficiently without argmin/argmax differentiation. We also prove the calibration of their excess risk in the case of linear-concave energies. We demonstrate our losses on multilabel classification and imitation learning tasks.

1 Introduction

Training a neural network usually involves finding its parameters by minimizing a loss function, which captures how well the network fits the data. A typical example is the multiclass logistic loss (a.k.a. cross-entropy loss) from logistic regression, which is the canonical loss associated with the softmax output layer and the categorical distribution. If we replace the softmax with another output layer, what loss function should we use instead? In generalized linear models [52, 48], which include logistic regression and Poisson regression as special cases, the negative log-likelihood gives a loss function associated with a link function, generalizing the softmax to other members of the exponential family [8]. More generally, **Fenchel-Young losses** [18] provide a generic way to construct a canonical convex loss function if the associated output layer can be written in a certain argmax form. Besides the aforementioned generalized linear models, models that fall in this family include sparsemax [47], the structured perceptron [27] and conditional random fields [43, 69]. However, the theory of Fenchel-Young losses is currently limited to argmax output layers that use a **bilinear** pairing.

To increase expressivity, energy-based models [44], a.k.a. **energy networks**, perform inference by optimizing an **energy function**, typically parametrized by a neural network. This leads to an inner optimization problem, which can capture potentially complex relationships between inputs and outputs. A similar approach is taken in SPENs (Structured Prediction Energy Networks) [11, 13], in which a continuous relaxation of the inner optimization problem is used, amenable to projected gradient descent or mirror descent. In input-convex neural networks [4], energy functions are restricted to be convex, so as to make the inner optimization problem easy to solve optimally. To learn the parameters of the energy function, the solution of the inner optimization problem is typically fed into a loss function, leading to an outer optimization problem. In order to solve that problem by stochastic gradient descent, the main challenge lies in computing the loss gradients. Indeed, when

using an arbitrary loss function, by the chain rule, computing gradients requires differentiating through the inner optimization problem solution, often referred to as **argmin** or **argmax differentiation**. This can be done by backpropagation through unrolled algorithm iterates [13] or implicit differentiation through the optimality conditions [4]. This issue can be circumvented by using a generalized perceptron loss [44] or a max margin loss [11]. These losses only require **max differentiation**, for which **envelope theorems** [49] can be used. However, these losses often fail to satisfy envelope theorem assumptions such as unicity of the solution and lack theoretical guarantees.

In this paper, we propose to extend the theory of Fenchel-Young losses [18], so as to create a canonical loss function associated with an energy network. Our proposal builds upon **generalized conjugate functions**, also known as **Fenchel-Moreau conjugates** [50, 62], allowing us to go beyond bilinear pairings and to support **general energy functions** such as neural networks. We introduce **regularization**, allowing us to obtain unicity of the solutions and to use envelope theorems to compute gradients, without argmin or argmax differentiation. We provide novel guarantees on the **calibration** of the excess risk, when our loss is used as a surrogate for another discrete loss. To sum up, we obtain a well-motivated loss construction for general energy networks. The rest of the paper is organized as follows.

- After providing some background (§2), we describe **regularized energy networks**, identify a classification of energy functions and give several possible examples that fall in this family (§3).
- We define **generalized conjugates**, an extension of Legendre-Fenchel conjugates, and state their properties (§4). We establish novel conditions for a generalized conjugate to be a smooth function.
- We then introduce **generalized Fenchel-Young losses** and show that they enjoy many of the same favorable properties as the regular Fenchel-Young losses (§5). On the theoretical side, we prove **calibration** of the excess risk for linear-concave energies and strongly-convex regularizers (§6).
- We demonstrate our losses on **multilabel classification** and **imitation learning** tasks (§7).

2 Background

Convex conjugates. Let $\mathcal{C} \subseteq \mathbb{R}^k$ be an output set. Let $\Omega: \mathcal{C} \rightarrow \mathbb{R}$ be a function, such that $\Omega(p) = \infty$ for all $p \notin \mathcal{C}$, i.e., $\text{dom}(\Omega) = \mathcal{C}$. Given $v \in \mathcal{V} \subseteq \mathbb{R}^k$, we define the convex conjugate [22] of Ω , also known as the Legendre-Fenchel transform of Ω , by

$$\Omega^*(v) := \max_{p \in \mathcal{C}} \langle v, p \rangle - \Omega(p). \quad (1)$$

The conjugate Ω^* is always convex, even if Ω is not. We define the corresponding argmax as

$$p_\Omega(v) := \operatorname{argmax}_{p \in \mathcal{C}} \langle v, p \rangle - \Omega(p) = \nabla \Omega^*(v). \quad (2)$$

The latter equality follows from Danskin's theorem [29, 15], under the assumption that Ω is strictly convex (otherwise, we obtain a subgradient). It is well-known that Ω^* is $\frac{1}{\gamma}$ -smooth w.r.t. the dual norm $\|\cdot\|_*$ if and only if Ω is γ -strongly convex w.r.t. the norm $\|\cdot\|$ [37, 39, 10, 75].

Fenchel-Young losses. Suppose that $v = g_\theta(x) \in \mathbb{R}^k$ are the logits / scores produced by a neural network g , where $x \in \mathcal{X} \subseteq \mathbb{R}^d$ and $\theta \in \Theta$ are the network's input features and parameters, respectively. What loss function should we use if we want to use (2) as output layer? The Fenchel-Young loss [18] generated by $\Omega: \mathcal{C} \rightarrow \mathbb{R}$ provides a natural solution. It is defined by

$$L_\Omega(v, y) := \Omega^*(v) + \Omega(y) - \langle v, y \rangle, \quad (3)$$

where $y \in \mathcal{Y} \subseteq \mathcal{C}$ is the ground-truth label. Earlier instances of this loss were independently proposed in the contexts of ranking [2] and multiclass classification [30]. Among many useful properties, this loss satisfies $L_\Omega(v, y) \geq 0$ and $L_\Omega(v, y) = 0 \Leftrightarrow y = p_\Omega(v)$ if Ω is strictly convex [18]. In that sense, it is the canonical loss associated with (2). Interestingly, many existing loss functions are recovered as special cases. For instance, if $\mathcal{C} = \mathbb{R}^k$ and $\Omega(p) = \frac{1}{2}\|p\|_2^2$, which is 1-strongly convex w.r.t. $\|\cdot\|_2$ over \mathbb{R}^k , then we obtain the self-dual, the identity mapping and the squared loss:

$$\Omega^*(v) = \frac{1}{2}\|v\|_2^2, \quad p_\Omega(v) = v, \quad L_\Omega(v, y) = \frac{1}{2}\|v - y\|_2^2.$$

If \mathcal{C} is the probability simplex $\Delta^k := \{p \in \mathbb{R}_+^k : \sum_{i=1}^k p_i = 1\}$ and $\Omega(p)$ is the scaled Shannon negentropy $\gamma \langle p, \log p \rangle$, which is γ -strongly convex w.r.t. $\|\cdot\|_1$ over Δ^k , then we obtain

$$\Omega^*(v) = \text{LSE}^\gamma(v), \quad p_\Omega(v) = \text{softmax}^\gamma(v), \quad L_\Omega(v, y) = \text{KL}(y, \text{softmax}^\gamma(v)),$$

where we used the log-sum-exp $\text{LSE}^\gamma(v) := \gamma \log(\sum_{i=1}^k \exp(v_i/\gamma))$, $\text{softmax}^\gamma(v) \propto \exp(v/\gamma)$, and the Kullback-Leibler divergence. More generally, when $\mathcal{C} = \text{conv}(\mathcal{Y})$, the convex hull of \mathcal{Y} , p_Ω corresponds to a projection, for which efficient algorithms exist in numerous cases [18, 16]. The calibration of the excess risk of Fenchel-Young losses when Ω is strongly convex was established in [56, 16]. However, the theory of Fenchel-Young losses is currently limited to the bilinear pairing $\langle v, p \rangle$, which restricts their expressivity and scope.

3 Regularized energy networks

Energy networks. Also known as energy-based models or EBMs [44], of which SPENs [11, 13] are a particular case, these networks compute predictions by solving an optimization problem of the form

$$p^\Phi(v) := \operatorname{argmax}_{p \in \mathcal{C}} \Phi(v, p),$$

where $v = g_\theta(x) \in \mathcal{V}$ is the energy network's input, $\Phi(v, p)$ is a scalar-valued energy function, and \mathcal{C} is an output set. Throughout this paper, we use the convention that higher energy indicates higher degree of compatibility between v and the prediction p . Since $\Phi(v, p)$ is a general energy function, we emphasize that v and p do not need to have the same dimensions, unlike with the bilinear pairing $\langle v, p \rangle$. Any neural network $g_\theta(x) \in \mathbb{R}^k$ can be written in energy network form, since $\operatorname{argmax}_{p \in \mathbb{R}^k} -\|g_\theta(x) - p\|_2^2 = g_\theta(x)$. The key advantage of energy networks, however, is their ability to capture complex interactions between inputs and outputs. This is especially useful in the structured prediction setting [6], where predictions are made of sub-parts, such as sequences being composed of individual elements.

Introducing regularization. In this paper, we compute predictions by solving

$$p_\Omega^\Phi(v) := \operatorname{argmax}_{p \in \mathcal{C}} \Phi(v, p) - \Omega(p), \quad (4)$$

where we further added a regularization function $\Omega: \mathcal{C} \rightarrow \mathbb{R}$. We call $\Phi(v, p) - \Omega(p)$ a **regularized energy function** and we call the corresponding model a **regularized energy network**. We obviously recover usual energy networks by setting Ω to the indicator function of the set \mathcal{C} , i.e., $\Omega(p) = 0$ if $p \in \mathcal{C}$, ∞ otherwise. While it is in principle possible to absorb Ω into Φ , keeping Ω explicit has several advantages: 1) it allows to introduce generalized conjugate functions and their properties (§4) 2) it allows to introduce generalized Fenchel-Young losses (§5), which mirror the original Fenchel-Young losses 3) we obtain closed forms for (4) in certain cases (Table 1).

Solving the maximization problem. The ability to solve the maximization problem in (4) efficiently depends on the properties of $\Phi(v, p)$ w.r.t. p . If $\Phi(v, p)$ is **linear** in p , for instance $\Phi(v, p) = \langle v, Up \rangle$, then we obtain $p_\Omega^\Phi(v) = p_\Omega(U^\top v)$. Thus, the computation of (4) reduces to (2), for which closed forms are often available. If $\Phi(v, p)$ is **concave** in p and \mathcal{C} is a convex set, then we can solve (4) in polynomial time using an iterative algorithm, such as projected gradient ascent. This is the most general energy class for which our loss and its gradient can be computed to arbitrary precision. If $\Phi(v, p)$ is nonconcave in p , then it is not possible to solve (4) in polynomial time in general, unless \mathcal{C} is a discrete set of small cardinality. In general, we emphasize that $\Phi(v, p)$ can be nonconvex in v , as is typical with neural networks, since the maximization problem in (4) is with respect to $p \in \mathcal{C}$. However, in the sequel of this paper, certain properties we establish require $\Phi(v, p)$ to be convex in v .

We now give examples of regularized energy networks in increasing order of expressivity / complexity.

Generalized linear models. As a warm up, we consider the case of **bilinear** energy $\Phi(v, p)$. For instance, for probabilistic classification with k classes, where the goal is to predict $p \in \Delta^k$ from $x \in \mathbb{R}^d$, we can use $\Phi(v, p) = \langle v, p \rangle$ with $v = Wx + b$, where $W \in \mathbb{R}^{k \times d}$ and $b \in \mathbb{R}^k$. In this case, we therefore obtain $p_\Omega^\Phi(v) = p_\Omega(Wx + b)$, recovering generalized linear models [52, 48], the structured perceptron [27] and conditional random fields [43, 69].

Rectifier and maxout networks. We now consider the case of **convex-linear** energy $\Phi(v, p)$. A first example is a rectifier network [33] with one hidden layer. Indeed, $\Phi(v, p) = \langle \sigma(v), Up \rangle$ is convex in v if σ is an element-wise, convex and non-decreasing activation function, such as the relu or softplus, and if U and p are non-negative. A second example is a maxout network [34] (i.e., a max-affine function) or its smoothed counterpart, the log-sum-exp network [24]. Indeed, $\Phi(v, p) = \sigma(v) \cdot p$, where σ is the max or log-sum-exp operator, is convex in v , if the scalar p is non-negative. In general, it is always possible to construct a **convex-concave** energy from a jointly-convex function using the Legendre-Fenchel transform in the first argument.

Table 1: Examples of regularized energy networks. The v and p columns indicate the property of the energy function $\Phi(v, p)$ in these variables. The $L_{\Omega}^{\Phi}(v, y)$ column indicates the property of the loss in v . The linear-quadratic energy uses $v = (A, b)$ where A is negative semi-definite and $\Omega(p) = \frac{\gamma}{2} \|p\|_2^2$.

	$\Phi(v, p)$	v	p	$L_{\Omega}^{\Phi}(v, y)$	$p_{\Omega}^{\Phi}(v)$
GLM	$\langle v, p \rangle$	linear	linear	convex	$p_{\Omega}(v)$
Linear-quadratic	$\frac{1}{2} \langle p, Ap \rangle + \langle p, b \rangle$	linear	quadratic	convex	$(\gamma I - A)^{-1} b$
Rectifier network	$\langle \text{relu}(v), Up \rangle$	convex	linear	DC	$p_{\Omega}(U^{\top} \text{relu}(v))$
Maxout network	$p \cdot \max(v)$	convex	linear	DC	$p_{\Omega}(\max(v))$
LSE network	$p \cdot \text{LSE}^{\gamma}(v)$	convex	linear	DC	$p_{\Omega}(\text{LSE}^{\gamma}(v))$
ICNN	$-\text{ICNN}(v, p)$	nonconvex	concave	nonconvex	no closed form
Probabilistic	$\sum_{y \in \mathcal{Y}} p(y) E(v, y)$	nonconvex	linear	nonconvex	$\frac{\exp(E(v, \cdot))}{\sum_{y' \in \mathcal{Y}} \exp(E(v, y'))}$
Arbitrary	$\Phi(v, p)$	nonconvex	nonconcave	nonconvex	no closed form

Input-convex neural networks. As an example of **nonconvex-concave** energy, we consider $\Phi(v, p) = -\text{ICNN}(v, p)$, where ICNN is an input-convex neural network [4], i.e., it is convex in p but can be nonconvex in v . If $\Omega(p)$ and \mathcal{C} are convex, then $\Phi(v, p)$ is concave in p and (4) can be solved in polynomial time using an iterative algorithm, such as projected gradient ascent.

Probabilistic energy networks. It is often desirable to define a conditional probability distribution

$$\mathbb{P}(Y = y | X = x) := \frac{\exp(E(v, y))}{\sum_{y' \in \mathcal{Y}} \exp(E(v, y'))}.$$

Such networks are typically trained using a cross-entropy loss (or equivalently, negative log-likelihood). Unfortunately, when \mathcal{Y} is large or infinite, this loss and its gradients are intractable to compute, due to the normalization constant. Therefore, recent research has been devoted to developing approximate training schemes [67, 35, 51]. Although this is not the focus of this paper, we point out that probabilistic energy networks can also be seen as regularized energy networks in the space of probability distributions $\mathcal{C} = \Delta^{|\mathcal{Y}|}$, if we set $\Phi(v, p) = \sum_{y \in \mathcal{Y}} p(y) E(v, y)$ and $\Omega(p) = \sum_{y \in \mathcal{Y}} p(y) \log p(y)$. While the resulting optimization problem is intractable in general, it does suggest possible approximation schemes, such as the use of Frank-Wolfe methods [12, 42, 55].

Existing loss functions for energy networks. Given a pair (x, y) and output $v = g_{\theta}(x)$, how do we measure the discrepancy between $p_{\Omega}^{\Phi}(v)$ and y ? One possibility [4, 13] is to use the composition of a differentiable loss $L: \mathcal{C} \times \mathcal{Y} \rightarrow \mathbb{R}_+$ with the argmax output, namely $(v, y) \mapsto L(p_{\Omega}^{\Phi}(v), y)$. However, computing the loss gradient w.r.t. v then requires to differentiate through $p_{\Omega}^{\Phi}(v)$, either through unrolling or implicit differentiation. This is particularly problematic when \mathcal{C} is a complicated convex set, as differentiating through a projection can be challenging. In contrast, our proposed loss completely circumvents this need and enjoys easy-to-compute gradients. For unregularized energy networks, a naive idea, called the energy loss, is to use $(v, y) \mapsto -\Phi(v, y)$. However, this loss only works well if Φ is a similarity measure and works poorly in general [44]. A better choice is the generalized perceptron loss $(v, y) \mapsto \max_{p \in \mathcal{C}} \Phi(v, p) - \Phi(v, y)$ [44]. Our loss can be seen as a principled generalization of this loss to regularized energy networks, with theoretical guarantees.

4 Generalized conjugates

In order to devise generalized Fenchel-Young losses, we build upon a generalization due to Moreau [50, Chapter 14] of the convex conjugate Ω^* . Denoted Ω^{Φ} , it replaces the bilinear pairing in (1) with a more general coupling Φ [62, Chapter 11, Section L]. In this section, we state their definition, properties, closed-form expressions and connection to the \mathcal{C} -transform in optimal transport.

Definition. Let $\Phi(v, p) \in \mathbb{R}$ be a coupling / energy function. The Φ -convex conjugate of $\Omega: \mathcal{C} \rightarrow \mathbb{R}$, also known as **Fenchel-Moreau conjugate**, is then defined by the value function

$$\Omega^{\Phi}(v) := \max_{p \in \mathcal{C}} \Phi(v, p) - \Omega(p). \quad (5)$$

The Φ -convex conjugate is an important tool in abstract convex analysis [66, 63]. Recently, it has been used to provide ‘‘Bellman-like’’ equations in stochastic dynamic programming [25] and to

provide tropical analogues of reproducing kernels [5]. We assume that the maximum is feasible for all $v \in \mathcal{V} = \mathbb{R}^d$, meaning that $\text{dom}(\Omega^\Phi) = \mathcal{V}$. We emphasize again that, unlike with (1), v and p do not need to have compatible dimensions. We denote the argmax solution corresponding to (5) by

$$p_\Omega^\Phi(v) := \operatorname{argmax}_{p \in \mathcal{C}} \Phi(v, p) - \Omega(p). \quad (6)$$

If a function $F(v)$ can be written as $F(v) = \Omega^\Phi(v)$ for some Ω , it is called Φ -convex (in analogy, a function $f(v)$ is convex and closed if and only if it can be written as $f(v) = \Omega^*(v)$ for some Ω).

Properties. Φ -convex conjugates enjoy many useful properties, some of them are natural extensions of the usual convex conjugate properties. Proofs are provided in Appendix B.1.

Proposition 1 (Properties of Φ -convex conjugates). *Let $\Omega: \mathcal{C} \rightarrow \mathbb{R}$ and $\Phi: \mathcal{V} \times \mathcal{C} \rightarrow \mathbb{R}$.*

1. **Generalized Fenchel-Young inequality:** for all $v \in \mathcal{V}$ and $p \in \mathcal{C}$,

$$\Omega^\Phi(v) + \Omega(p) - \Phi(v, p) \geq 0.$$

2. **Convexity:** If $\Phi(v, p)$ is convex in v , then $\Omega^\Phi(v)$ is convex (even if $\Omega(p)$ is nonconvex).

3. **Order reversing:** if $\Omega(p) \leq \Lambda(p)$ for all $p \in \mathcal{C}$, then $\Omega^\Phi(v) \geq \Lambda^\Phi(v)$ for all $v \in \mathcal{V}$.

4. **Continuity:** Ω^Φ shares the same continuity modulus as Φ .

5. **Gradient (envelope theorem):** Under assumptions (see paragraph below), we have $\nabla \Omega^\Phi(v) = \nabla_1 \Phi(v, p_\Omega^\Phi(v))$, where ∇_1 denotes the gradient in the first argument.

6. **Smoothness:** If \mathcal{C} is a compact convex set, $\Phi(v, p)$ is β -smooth in (v, p) , concave in p and $\Omega(p)$ is γ -strongly convex in p , then $\Omega^\Phi(v)$ is $(\beta + \beta^2/\gamma)$ -smooth and $p_\Omega^\Phi(v)$ is β/γ -Lipschitz.

The condition on Φ and Ω in item 6 for Ω^Φ to be a smooth function (i.e., with Lipschitz-continuous gradients) is a novel result and will play a crucial role for establishing calibration guarantees in §6.

Assumptions for envelope theorems. The expression in item 5 allows to compute $\nabla \Omega^\Phi(v)$ without *argmax* differentiation. It is based on envelope theorems, which can be used for *max* differentiation. Indeed, we have $\Omega^\Phi(v) = \max_{p \in \mathcal{C}} F(v, p)$, where $F(v, p) := \Phi(v, p) - \Omega(p)$. We assume that the maximum is unique and \mathcal{C} is a compact. If $F(v, p)$ is convex in v , we apply Danskin's theorem [29] [15, Proposition B.25]. Without convexity assumption in v , if $F(v, p)$ is continuously differentiable in v for all $p \in \mathcal{C}$, $\nabla_1 F$ is continuous, we apply [62, Theorem 10.31]. When we do not compute the exact solution of (6), we only obtain an approximation of the gradient $\nabla \Omega^\Phi(v)$; see [1] for approximation guarantees. For other envelope theorem usecases in machine learning, see, e.g., [21].

Closed forms. While (5) and (6) may need to be solved numerically in general, they enjoy closed-form expressions in simple cases. Proofs are provided in Appendix B.2.

Proposition 2 (Closed-form expressions). *Let $\Omega: \mathcal{C} \rightarrow \mathbb{R}$ and $\Phi: \mathcal{V} \times \mathcal{C} \rightarrow \mathbb{R}$.*

1. **Bilinear coupling:** If $\Phi(v, p) = \langle v, Up \rangle$, then $\Omega^\Phi(v) = \Omega^*(U^\top v)$ and $p_\Omega^\Phi(v) = p_\Omega(U^\top v)$.

2. **Linear-quadratic coupling:** If $\mathcal{C} = \mathbb{R}^k$, $\Omega(p) = \frac{\gamma}{2} \|p\|_2^2$ and $\Phi(v, p) = \frac{1}{2} \langle p, Ap \rangle + \langle p, b \rangle$, where $v = (A, b)$ and A is such that $(\gamma I - A)$ is positive definite, we obtain

$$\Omega^\Phi(v) = \frac{1}{2} \langle b, (\gamma I - A)^{-1} b \rangle \quad \text{and} \quad p_\Omega^\Phi(v) = (\gamma I - A)^{-1} b. \quad (7)$$

3. **Metric coupling:** If $\mathcal{V} = \mathcal{C}$, $\Phi = -C$ where $C(v, p)$ is a metric and Ω is M -Lipschitz with $M \leq 1$, then $\Omega^\Phi = -\Omega$ and $p_\Omega^\Phi(v) = v$.

Relation with the C -transform. Given a cost function C , we may define a min counterpart of (5),

$$\Lambda_C(v) := \min_{p \in \mathcal{C}} C(v, p) - \Lambda(p). \quad (8)$$

In the optimal transport literature, this is known as the C -transform of Λ [65, 60]. When C is bilinear, this recovers the notion of concave conjugate [20]. When $C(v, p) = c(v - p)$ for some c , this recovers the infimal convolution [20]. It is easy to check that $\Omega = -\Lambda \Leftrightarrow \Omega^\Phi = -\Lambda_C$ with $C = -\Phi$. Thus, Ω^Φ and Λ_C are the natural extensions of convex and concave conjugates, respectively. We opt for the former in this paper to closely mirror the usual convex conjugates and Fenchel-Young losses.

5 Generalized Fenchel-Young losses

Definition. The generalized Fenchel-Young inequality in Proposition 1 leads us to propose the generalized Fenchel-Young loss, the natural extension of (3) to Φ -convex conjugates:

$$L_{\Omega}^{\Phi}(v, y) := \Omega^{\Phi}(v) + \Omega(y) - \Phi(v, y). \quad (9)$$

We also have the relationships $L_{\Omega}^{\Phi}(v, y) = F(v, y) - F(v, p_{\Omega}^{\Phi}(v))$, where $F(v, p) = \Omega(p) - \Phi(v, p)$, and $p_{\Omega}^{\Phi}(v) = \operatorname{argmin}_{p \in \mathcal{C}} L_{\Omega}^{\Phi}(v, p)$. We now give an intuitive geometric interpretation of (9).

Geometric interpretation. For the sake of illustration, let us set $\Omega(p) = \frac{\gamma}{2}p^2$ and $\Phi(v, p) = \frac{1}{2}ap^2 + bp$. From (1), the line $p \mapsto \langle u, p \rangle - \Omega^*(u)$, where $\Omega^*(u) = \frac{1}{2\gamma}u^2$, is the tightest linear lower bound on $\Omega(p)$, here depicted with $u = 0.6$. It is the tangent of $\Omega(p)$ at $p = p_{\Omega}(u)$. Similarly, from (5), $p \mapsto \Phi(v, p) - \Omega^{\Phi}(v)$ is the tightest lower-bound of this form given v . It touches $\Omega(p)$ at $p = p_{\Omega}^{\Phi}(v)$, here depicted with $v = (a, b)$, $a = -1$ and $b = 1$, for which we have the closed form $\Omega^{\Phi}(v) = \frac{1}{2}(\gamma - a)^{-1}b^2$ (Proposition 2). The generalized Fenchel-Young loss (9) is then the **gap** between $\Omega(p)$ and $\Phi(v, p) - \Omega^{\Phi}(v)$, evaluated at the ground-truth $p = y$. The goal of training is to adjust the parameters θ of a network $v = g_{\theta}(x)$ so as to minimize this gap averaged over all (x, y) training pairs.

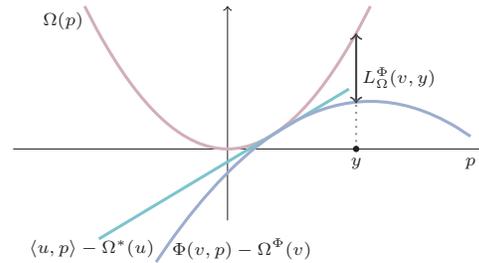


Figure 1: Geometric interpretation with $\Omega(p) = \frac{\gamma}{2}p^2$ and $\Phi(v, p) = \frac{1}{2}ap^2 + bp$.

Properties. Generalized Fenchel-Young losses enjoy many desirable properties, as we now show.

Proposition 3 (Properties of generalized F-Y losses). *Let $\Omega: \mathcal{C} \rightarrow \mathbb{R}$ and $\Phi: \mathcal{V} \times \mathcal{C} \rightarrow \mathbb{R}$.*

1. **Non-negativity:** $L_{\Omega}^{\Phi}(v, p) \geq 0$ for all $v \in \mathcal{V}$ and $p \in \mathcal{C}$.
2. **Zero loss:** if the maximum in (5) exists and is unique, $L_{\Omega}^{\Phi}(v, y) = 0 \Leftrightarrow y = p_{\Omega}^{\Phi}(v)$.
3. **Gradient:** $\nabla_1 L_{\Omega}^{\Phi}(v, p) = \nabla \Omega^{\Phi}(v) - \nabla_1 \Phi(v, p)$.
4. **Difference of convex (DC):** if $\Phi(v, p)$ is convex in v , then $L_{\Omega}^{\Phi}(v, p)$ is a difference of convex functions in v : $\Omega^{\Phi}(v)$ and $\Phi(v, p)$. If $\Phi(v, p)$ is linear in v , then $L_{\Omega}^{\Phi}(v, p)$ is convex in v .
5. **Smaller output set, smaller loss.** If $\mathcal{C}' \subseteq \mathcal{C}$ and Ω' is the restriction of Ω to \mathcal{C}' , then $L_{\Omega'}^{\Phi}(v, p) \leq L_{\Omega}^{\Phi}(v, p)$ for all $v \in \mathcal{V}$ and $p \in \mathcal{C}'$.
6. **Quadratic lower-bound.** If $\Phi(v, p) - \Omega(p)$ is γ -strongly concave in p w.r.t. $\|\cdot\|$ over \mathcal{C} and \mathcal{C} is a closed convex set, then for all $v \in \mathcal{V}$ and $p \in \mathcal{C}$

$$\frac{\gamma}{2} \|p - p_{\Omega}^{\Phi}(v)\|^2 \leq L_{\Omega}^{\Phi}(v, p). \quad (10)$$

7. **Upper-bounds.** If $\Phi(v, p) - \Omega(p)$ is α -Lipschitz in p w.r.t. $\|\cdot\|$ over \mathcal{C} , then $L_{\Omega}^{\Phi}(v, p) \leq \alpha \|p - p_{\Omega}^{\Phi}(v)\|$. If $\Phi(v, p)$ is concave in p , then $L_{\Omega}^{\Phi}(v, p) \leq L_{\Omega}(\nabla_2 \Phi(v, p), p)$.

Proofs are given in Appendix B.3. The **non-negativity** is a good property for a loss function. The **zero-loss** property $L_{\Omega}^{\Phi}(v, y) = 0 \Leftrightarrow y = p_{\Omega}^{\Phi}(v)$, which is true for instance if $p \mapsto \Phi(v, p) - \Omega(p)$ is strictly concave, is key as it allows to use $p_{\Omega}^{\Phi}(v)$ defined in (6) as the (implicit) output layer associated with an energy network. The **gradient** $\nabla_1 L_{\Omega}^{\Phi}(v, y)$ does not require to differentiate through the argmax problem in (6) needed for computing $p_{\Omega}^{\Phi}(v)$. Typically, differentiating through an argmax or argmin, as is done in input-convex neural networks [4], requires either unrolling or implicit differentiation [36, 14, 41, 19, 17] and is therefore more costly.

The fact that $L_{\Omega}^{\Phi}(v, p)$ is a **difference of convex (DC)** functions in v when $\Phi(v, p)$ is convex in v suggests that we can use DC programming techniques, such as the convex-concave procedure [72], for training such energy networks. We leave the investigation of this observation to future work. The **“smaller output set, smaller loss”** property means that we can achieve the smallest loss by choosing the smallest set \mathcal{C} in (6) such that $\mathcal{Y} \subseteq \mathcal{C}$. The smallest such convex set is the convex hull of \mathcal{Y} , also

known as marginal polytope [71] when $\mathcal{Y} \subseteq \{0, 1\}^k$. The **quadratic lower-bound** relates our loss to using $p_{\Omega}^{\Phi}(v)$ within a squared norm loss. The **upper-bounds** relate our loss to using $p_{\Omega}^{\Phi}(v)$ within a norm loss and to using a regular Fenchel-Young loss with a linearized energy.

If Ω is the indicator function of \mathcal{C} , i.e., $\Omega(p) = 0$ if $p \in \mathcal{C}$, ∞ otherwise, then it can be checked that we recover the “generalized perceptron” loss [45, 44] as a special case of (9). Proposition 3 therefore provides new properties to understand and analyze this loss.

Regular Fenchel-Young losses are closely related to Bregman divergences [18]. In Appendix A, we build a generalized notion of Bregman divergence using generalized conjugates.

Training. To train regularized energy networks with our framework, the user should choose an energy $\Phi(v, p)$, a regularization $\Omega(p)$, an output set \mathcal{C} (from Proposition 3 item 5, the smaller this set the better) and the model $v = g_{\theta}(x)$ with input x and parameters θ . Given a set of input-output pairs $(x_1, y_1), \dots, (x_n, y_n) \in \mathcal{X} \times \mathcal{Y}$, where $\mathcal{Y} \subseteq \mathcal{C}$, we can find the parameters θ by minimizing the empirical risk objective regularized by $R: \Theta \rightarrow \mathbb{R}$,

$$\hat{\theta} = \operatorname{argmin}_{\theta \in \Theta} \frac{1}{n} \sum_{i=1}^n L_{\Omega}^{\Phi}(g_{\theta}(x_i), y_i) + R(\theta). \quad (11)$$

Thanks to the easy-to-compute gradients of L_{Ω}^{Φ} , we can easily solve (11) using any (stochastic) solver.

6 Calibration guarantees

Many times, notably for differentiability reasons, the loss used at training time, here our generalized Fenchel-Young loss $L_{\Omega}^{\Phi}(v, y)$, is used as a surrogate / proxy for a different (potentially discrete loss) $L: \mathcal{Y} \times \mathcal{Y} \rightarrow \mathbb{R}$, used at test time. Calibration guarantees [73, 74, 9, 68] ensure that minimizing the excess of risk of the train loss will also minimize that of the test loss (a.k.a. target loss). We study in this section such guarantees, assuming that L satisfies an affine decomposition property [26, 16]:

$$L(\hat{y}, y) = \langle \varphi(\hat{y}), V\varphi(y) + b \rangle + c(y), \quad (12)$$

where $\varphi(y) \mapsto V\varphi(y) + b$ is an affine map, $\varphi(y)$ is a label embedding and $c(y)$ is any function that depends only on y . Numerous losses can be written in this form. Examples include the zero-one, Hamming, NDCG and precision at k losses [56, 16]. Inference in this setting works in two steps. First, we compute a “soft” (continuous) prediction $p = p_{\Omega}^{\Phi}(v) \in \mathcal{C}$. Second, we compute a “hard” (discrete) prediction by a decoding / rounding from \mathcal{C} to \mathcal{Y} , calibrated for the loss L :

$$y_L(p) := \operatorname{argmin}_{\hat{y} \in \mathcal{Y}} L(\hat{y}, p) = \operatorname{argmin}_{\hat{y} \in \mathcal{Y}} \langle \varphi(\hat{y}), V\varphi(p) + b \rangle. \quad (13)$$

The target risk of $f: \mathcal{X} \rightarrow \mathcal{Y}$ and the surrogate risk of $g: \mathcal{X} \rightarrow \mathcal{Y}$ are defined by

$$\mathcal{L}(f) := \mathbb{E}_{(X, Y) \sim \rho} L(f(X), Y) \quad \text{and} \quad \mathcal{L}_{\Omega}^{\Phi}(g) := \mathbb{E}_{(X, Y) \sim \rho} L_{\Omega}^{\Phi}(g(X), Y),$$

where ρ is a typically unknown distribution over $\mathcal{X} \times \mathcal{Y}$. The Bayes predictors are defined by $f^* := \operatorname{argmin}_{f: \mathcal{X} \rightarrow \mathcal{Y}} \mathcal{L}(f)$ and $g^* := \operatorname{argmin}_{g: \mathcal{X} \rightarrow \mathcal{Y}} \mathcal{L}_{\Omega}^{\Phi}(g)$. We now establish calibration of the surrogate excess risk, under the assumption that the energy $\Phi(v, p)$ is **linear-concave**, i.e., it can be written as $\Phi(v, p) = \langle v, \varphi(p) \rangle$, for some function φ .

Proposition 4. *Calibration of target and surrogate excess risks*

Assume $L_{\Omega}^{\Phi}(v, y)$ is M -smooth in v w.r.t. the dual norm $\|\cdot\|_*$, \mathcal{C} is a compact convex set such that $\mathcal{Y} \subseteq \mathcal{C}$ and $\Phi(v, p) = \langle v, \varphi(p) \rangle$. Let $\sigma := \sup_{y \in \mathcal{Y}} \|V^{\top} \varphi(y)\|_*$. Then, the generalized Fenchel-Young loss (9) is calibrated with the target loss (12) with decoder $d = y_L \circ p_{\Omega}^{\Phi}$:

$$\forall g: \mathcal{X} \rightarrow \mathcal{Y} \quad \frac{(\mathcal{L}(d \circ g) - \mathcal{L}(f^*))^2}{8\sigma^2 M} \leq \mathcal{L}_{\Omega}^{\Phi}(g) - \mathcal{L}_{\Omega}^{\Phi}(g^*),$$

where $p_{\Omega}^{\Phi}: \mathcal{V} \rightarrow \mathcal{C}$ is defined in (6) and $y_L: \mathcal{C} \rightarrow \mathcal{Y}$ is defined in (13).

The proof is in Appendix B.4. Proposition 1 item 6 shows that the smoothness of Φ and strong convexity of Ω ensure the smoothness of L_{Ω}^{Φ} . Calibration also implies Fisher consistency, namely

Table 2: Multilabel classification results using various energies (test accuracy in %).

Energy	yeast	scene	mediamill	birds	emotions	cal500
Unary (linear)	79.76	89.14	96.84	86.47	78.22	85.67
Unary (rectifier network)	80.03	91.35	96.91	91.74	79.79	86.25
Pairwise	80.19	91.58	96.95	91.55	80.56	85.73
SPEN	79.99	91.24	96.68	91.41	79.35	86.25
Input-concave SPEN	80.00	90.64	96.95	91.77	79.73	86.35

$\mathcal{L}_\Omega^\Phi(g) = \mathcal{L}_\Omega^\Phi(g^*) \Rightarrow \mathcal{L}(d \circ g) = \mathcal{L}(f^*)$, when using the decoder $d = y_L \circ p_\Omega^\Phi$. The existing proof technique for the calibration of regular Fenchel-Young losses [56, 16] assumes a bilinear pairing and a loss of the form $L_{\Omega_\varphi}(u, \varphi(y)) = \Omega_\varphi^*(u) + \Omega_\varphi(\varphi(y)) - \langle u, \varphi(y) \rangle$, where Ω_φ is a strongly-convex regularizer w.r.t. $\varphi(y)$. Our novel proof technique is more general, as it works with any linear-concave energy and the regularizer Ω is w.r.t. y , not $\varphi(y)$. Unlike the existing proof, our proof is valid for the pairwise model we present in the next section.

7 Experiments

7.1 Multilabel classification

We study in this section the application of generalized Fenchel-Young losses to multi-label classification, setting $\mathcal{Y} = \{0, 1\}^k$ and $\mathcal{C} = [0, 1]^k$, where k is the number of labels. When the loss L in (12) is the Hamming loss (1 - accuracy), our loss is calibrated for L and the decoding (13) is just $\hat{y}_j = 1$ if $p_j > 0.5$ else 0. We therefore report our empirical results using the accuracy metric.

Unary model. We consider a neural network $u = g_\theta(x) \in \mathbb{R}^k$, assigning a score u_j to each label $j \in [k]$. With the bilinear pairing $\Phi(u, p) = \langle u, p \rangle$, we get

$$p_\Omega^\Phi(u) = \operatorname{argmax}_{p \in [0, 1]^k} \langle u, p \rangle - \Omega(p) = p_\Omega(u). \quad (14)$$

That is, (14) is just a normal neural network with p_Ω as output layer. When $\Omega(p) = \Omega_1(p) + \Omega_1(1-p)$, where $\Omega_1(p) := \langle p, \log p \rangle$ is Shannon's negentropy, we get $p_\Omega(v) = \operatorname{sigmoid}(v) := 1/(1 + \exp(-v))$ and (9) is just the usual binary logistic / cross-entropy loss. When $\Omega(p) = \Omega_2(p) + \Omega_2(1-p)$, where $\Omega_2(p) := \frac{1}{2} \langle p, 1-p \rangle$ is Gini's negentropy, we get a sparse sigmoid and the binary sparsemax loss [18, §6.2]. Because $\mathcal{C} = [0, 1]^k$ is the convex hull of $\mathcal{Y} = \{0, 1\}^k$, (14) can be interpreted as a **marginal probability** [71]. Indeed, there exists a probability distribution $\mathbb{P}(Y|X)$ over $Y \in \mathcal{Y}$ such that

$$[p_\Omega^\Phi(u)]_j = \mathbb{P}(Y_j = 1 | X = x).$$

Pairwise model. We now additionally use a network $U = h_\theta(x) \in \mathbb{R}^{k \times k}$, assigning a score $U_{i,j}$ to the pairwise interaction between labels i and j . With $v = (u, U)$ and the **linear-quadratic** coupling $\Phi(v, p) = \langle u, p \rangle + \frac{1}{2} \langle p, Up \rangle$, we get

$$p_\Omega^\Phi(v) = \operatorname{argmax}_{p \in [0, 1]^k} \langle u, p \rangle + \frac{1}{2} \langle p, Up \rangle - \Omega(p). \quad (15)$$

If U is negative semi-definite, i.e., $U = -AA^\top$ for some matrix $A \in \mathbb{R}^{k \times m}$, then the problem is concave in p and can be solved optimally. Moreover, since $\Phi(v, p)$ is linear-concave, the calibration guarantees in §6 hold and $L_\Omega^\Phi(v, y)$ is convex in v . In our experiments, we use a matrix A of rank 1 (cf. Appendix C). Unlike (7), (15) does not enjoy a closed form due to the constraints. We solve it by coordinate ascent: if Ω is quadratic, as is the case with Gini's negentropy, the coordinate-wise updates can be computed in closed-form. Again, (15) can be interpreted as a marginal probability. In contrast to (15), marginal inference in the closely related **Ising model** is known to be #P-hard [32].

SPEN model. Following SPENs [11, Eq. 4 and 5], we also tried the energy $\Phi(v, p) = \langle u, p \rangle - \Psi(w, p)$, where $v = (u, w)$, $u = g_\theta(x)$ and w are the weights of the "prior network" Ψ (independent of x). We also tried a variant where Ψ is made convex in p , making $\Phi(v, p)$ concave in p . In both cases, we compute $p_\Omega^\Phi(v)$ by solving (6) using projected gradient ascent with backtracking linesearch.

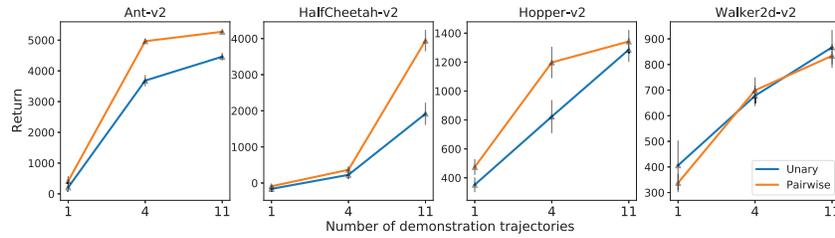


Figure 2: Average performance (higher is better) and standard deviation over 10 seeds.

Experimental setup. We perform experiments on 6 publicly-available datasets, see Appendix C. We use the train-test split from the dataset when provided. When not, we use 80% for training data and 20% for test data. For all models, we solve the outer problem (11) using ADAM. We set $\Omega(p)$ to the Gini negentropy. We hold out 25% of the training data for hyperparameter validation purposes. We set $R(\theta)$ in (11) to $\frac{\lambda}{2} \|\theta\|_2^2$. For the regularization hyper-parameter λ , we search 5 log-spaced values between 10^{-4} and 10^1 . For the learning rate parameter of ADAM, we search 10 log-spaced values between 10^{-5} and 10^{-1} . Once we selected the best hyperparameters, we refit the model on the entire training set. We average results over 3 runs with different seeds.

Results. Table 2 shows a model comparison. We observe improvements with the pairwise model on 4 out of 6 datasets, confirming that our losses are able to learn useful models. Input-concavity helps improve SPENs in 4 out of 6 datasets. Table 4 confirms that using the envelope theorem for computing gradients works comparably to (if not better than) the implicit function theorem. Table 5 shows that our losses outperform the energy, the cross-entropy and the generalized perceptron losses.

7.2 Imitation learning

In this section, we study the application of generalized Fenchel-Young losses to imitation learning. This setting consists in learning a policy $\pi : \mathcal{X} \mapsto \mathcal{Y}$, a mapping from states \mathcal{X} to actions \mathcal{Y} , from a fixed dataset of expert demonstrations $(x_1, y_1), \dots, (x_n, y_n) \in \mathcal{X} \times \mathcal{Y}$. In particular, we only consider a Behavior Cloning approach [61], which essentially reduces imitation learning to supervised learning (as opposed to inverse RL methods [64, 54]). The learned policy π is evaluated on its performance, which is the expected sum of rewards of the environment [70].

Experimental setup. We consider four MuJoCo Gym locomotion environments [23] together with the demonstrations provided by Orsini et al. [58]; see Appendix C.2 for details. We evaluate the learned policy π for different number of demonstration trajectories: 1, 4 and 11, consistently with [38, 40, 28]. The action space in the demonstrations is included in $[-1, 1]^k$, where the dimensionality of the action space k corresponds to the torques of the actuators. We scale the action space to the hypercube $\mathcal{Y} = \mathcal{C} = [0, 1]^k$ at learning time and scale it back to the original action space at inference time. Similarly to the multilabel classification setup, we evaluate the unary and pairwise models, with the only difference that we use two hidden layers instead of one, since it leads to significantly better performance. We specify the hyperparameter selection procedure in Appendix C.2.

Results. We run the best hyperparameters over 10 seeds and report final performance over 100 evaluation episodes. Figure 2 shows a clear improvement of the pairwise model over the unary model for 3 out of 4 tasks. Contrary to the unary model, the pairwise model enables to capture the interdependence between the different torques of the action space, translating into better performance.

8 Conclusion

Building upon generalized conjugate functions, we proposed generalized Fenchel-Young losses, a natural loss construction for learning energy networks and studied its properties. Thanks to conditions on the energy Φ and the regularizer Ω ensuring the loss smoothness, we established calibration guarantees for the case of linear-concave energies, a more general result than the existing analysis, restricted to bilinear energies. We demonstrated the effectiveness of our losses on multilabel classification and imitation learning tasks. We hope that this paper will help popularize generalized conjugates as a powerful tool for machine learning and optimization.

Acknowledgments

MB thanks Gabriel Peyré for numerous discussions on C -transforms and Clarice Poon for discussions on envelope theorems, as well as Vlad Niculae and André Martins for many fruitful discussions.

References

- [1] P. Ablin, G. Peyré, and T. Moreau. Super-efficiency of automatic differentiation for functions defined as a minimum. In *Proc. of ICML*, pages 32–41. PMLR, 2020.
- [2] S. Acharyya. Learning to rank in supervised and unsupervised settings using convexity and monotonicity. 2013.
- [3] S. Amari. *Information Geometry and Its Applications*. Springer, 2016.
- [4] B. Amos, L. Xu, and J. Z. Kolter. Input convex neural networks. In *ICML*, pages 146–155. PMLR, 2017.
- [5] P.-C. Aubin-Frankowski and S. Gaubert. The tropical analogues of reproducing kernels. *arXiv preprint arXiv:2202.11410*, 2022.
- [6] G. BakIr, T. Hofmann, B. Schölkopf, A. J. Smola, and B. Taskar. *Predicting structured data*. MIT press, 2007.
- [7] A. Banerjee, S. Merugu, I. S. Dhillon, J. Ghosh, and J. Lafferty. Clustering with bregman divergences. *Journal of machine learning research*, 6(10), 2005.
- [8] O. Barndorff-Nielsen. *Information and Exponential Families: In Statistical Theory*. John Wiley & Sons, 1978.
- [9] P. L. Bartlett, M. I. Jordan, and J. D. McAuliffe. Convexity, classification, and risk bounds. *Journal of the American Statistical Association*, 101(473):138–156, 2006.
- [10] A. Beck. *First-order methods in optimization*. SIAM, 2017.
- [11] D. Belanger and A. McCallum. Structured prediction energy networks. In *International Conference on Machine Learning*, pages 983–992. PMLR, 2016.
- [12] D. Belanger, D. Sheldon, and A. McCallum. **Marginal inference in MRFs using Frank-Wolfe**. In *NeurIPS Workshop on Greedy Opt., FW and Friends*, 2013.
- [13] D. Belanger, B. Yang, and A. McCallum. End-to-end learning for structured prediction energy networks. In *International Conference on Machine Learning*, pages 429–439. PMLR, 2017.
- [14] B. M. Bell and J. V. Burke. Algorithmic differentiation of implicit functions and optimal values. In *Advances in Automatic Differentiation*, pages 67–77. Springer, 2008.
- [15] D. P. Bertsekas. Nonlinear programming. *Journal of the Operational Research Society*, 48(3):334–334, 1997.
- [16] M. Blondel. Structured prediction with projection oracles. *Advances in neural information processing systems*, 32, 2019.
- [17] M. Blondel, Q. Berthet, M. Cuturi, R. Frostig, S. Hoyer, F. Llinares-López, F. Pedregosa, and J.-P. Vert. Efficient and modular implicit differentiation. *arXiv preprint arXiv:2105.15183*, 2021.
- [18] M. Blondel, A. F. Martins, and V. Niculae. Learning with fenchel-young losses. *JMLR*, 21(35):1–69, 2020.
- [19] J. F. Bonnans and A. Shapiro. *Perturbation analysis of optimization problems*. Springer Science & Business Media, 2013.
- [20] J. Borwein and A. Lewis. *Convex Analysis*. 2006.

- [21] L. Bottou, M. Arjovsky, D. Lopez-Paz, and M. Oquab. Geometrical insights for implicit generative modeling. pages 229–268. 2018.
- [22] S. Boyd, S. P. Boyd, and L. Vandenberghe. *Convex optimization*. Cambridge university press, 2004.
- [23] G. Brockman, V. Cheung, L. Pettersson, J. Schneider, J. Schulman, J. Tang, and W. Zaremba. Openai gym, 2016.
- [24] G. C. Calafiore, S. Gaubert, and C. Possieri. Log-sum-exp neural networks and posynomial models for convex and log-log-convex data. *IEEE transactions on neural networks and learning systems*, 31(3):827–838, 2019.
- [25] J.-P. Chancelier and M. De Lara. Fenchel-moreau conjugation inequalities with three couplings and application to stochastic bellman equation. *arXiv preprint arXiv:1804.03034*, 2018.
- [26] C. Ciliberto, L. Rosasco, and A. Rudi. A consistent regularization approach for structured prediction. In *Proc. of NeurIPS*, pages 4412–4420, 2016.
- [27] M. Collins. Discriminative training methods for Hidden Markov Models: Theory and experiments with perceptron algorithms. In *Proc. of EMNLP*, 2002.
- [28] R. Dadashi, L. Hussenot, M. Geist, and O. Pietquin. Primal wasserstein imitation learning. *International Conference on Learning Representations*, 2021.
- [29] J. M. Danskin. *The theory of max-min and its application to weapons allocation problems*. Springer Science & Business Media, 1967.
- [30] J. C. Duchi, K. Khosravi, and F. Ruan. Multiclass classification, information, divergence, and surrogate risk. *The Annals of Statistics*, 46(6B):3246–3275, 2018.
- [31] S. Ghadimi and M. Wang. Approximation methods for bilevel programming. *arXiv preprint arXiv:1802.02246*, 2018.
- [32] A. Globerson, T. Roughgarden, D. Sontag, and C. Yildirim. How hard is inference for structured prediction? In *International Conference on Machine Learning*, pages 2181–2190. PMLR, 2015.
- [33] X. Glorot, A. Bordes, and Y. Bengio. Deep sparse rectifier neural networks. In *AISTATS*, pages 315–323. JMLR Workshop and Conference Proceedings, 2011.
- [34] I. Goodfellow, D. Warde-Farley, M. Mirza, A. Courville, and Y. Bengio. Maxout networks. In *ICML*, pages 1319–1327. PMLR, 2013.
- [35] W. Grathwohl, J. Kelly, M. Hashemi, M. Norouzi, K. Swersky, and D. Duvenaud. No mcmc for me: Amortized samplers for fast and stable training of energy-based models. 2021.
- [36] A. Griewank and A. Walther. *Evaluating derivatives: principles and techniques of algorithmic differentiation*. SIAM, 2008.
- [37] J.-B. Hiriart-Urruty and C. Lemaréchal. *Convex analysis and minimization algorithms II*, volume 305. Springer science & business media, 1993.
- [38] J. Ho and S. Ermon. Generative adversarial imitation learning. In *Advances in Neural Information Processing Systems*, 2016.
- [39] S. Kakade, S. Shalev-Shwartz, A. Tewari, et al. On the duality of strong convexity and strong smoothness: Learning applications and matrix regularization. *Tech report*, 2(1):35, 2009.
- [40] I. Kostrikov, K. K. Agrawal, D. Dwibedi, S. Levine, and J. Tompson. Discriminator-actor-critic: Addressing sample inefficiency and reward bias in adversarial imitation learning. *International Conference on Learning Representations*, 2019.
- [41] S. G. Krantz and H. R. Parks. *The implicit function theorem: history, theory, and applications*. Springer Science & Business Media, 2012.

- [42] R. G. Krishnan, S. Lacoste-Julien, and D. Sontag. [Barrier Frank-Wolfe for marginal inference](#). In *Proc. of NeurIPS*, 2015.
- [43] J. D. Lafferty, A. McCallum, and F. C. Pereira. [Conditional Random Fields: Probabilistic models for segmenting and labeling sequence data](#). In *Proc. of ICML*, 2001.
- [44] Y. LeCun, S. Chopra, R. Hadsell, M. Ranzato, and F. Huang. A tutorial on energy-based learning. *Predicting structured data*, 1(0), 2006.
- [45] Y. LeCun and F. J. Huang. Loss functions for discriminative training of energy-based models. In *International Workshop on Artificial Intelligence and Statistics*, pages 206–213. PMLR, 2005.
- [46] T. Lin, C. Jin, and M. Jordan. On gradient descent ascent for nonconvex-concave minimax problems. In *International Conference on Machine Learning*, pages 6083–6093. PMLR, 2020.
- [47] A. F. Martins and R. F. Astudillo. [From softmax to sparsemax: A sparse model of attention and multi-label classification](#). In *Proc. of ICML*, 2016.
- [48] P. McCullagh and J. A. Nelder. [Generalized Linear Models](#), volume 37. CRC press, 1989.
- [49] P. Milgrom and I. Segal. Envelope theorems for arbitrary choice sets. *Econometrica*, 70(2):583–601, 2002.
- [50] J.-J. Moreau. Fonctionnelles convexes. *Séminaire Jean Leray*, (2):1–108, 1966.
- [51] C. Nash and C. Durkan. Autoregressive energy machines. In *International Conference on Machine Learning*, pages 1735–1744. PMLR, 2019.
- [52] J. A. Nelder and R. J. Baker. [Generalized Linear Models](#). Wiley Online Library, 1972.
- [53] Y. Nesterov. *Introductory lectures on convex optimization: A basic course*, volume 87. Springer Science & Business Media, 2003.
- [54] A. Y. Ng, S. J. Russell, et al. Algorithms for inverse reinforcement learning. In *International Conference on Machine Learning*, 2000.
- [55] V. Niculae, A. F. Martins, M. Blondel, and C. Cardie. [SparseMAP: Differentiable sparse structured inference](#). In *Proc. of ICML*, 2018.
- [56] A. Nowak-Vila, F. Bach, and A. Rudi. A general theory for structured prediction with smooth convex surrogates. *arXiv preprint arXiv:1902.01958*, 2019.
- [57] F. Orabona. A modern introduction to online learning. *arXiv preprint arXiv:1912.13213*, 2019.
- [58] M. Orsini, A. Raichuk, L. Hussenot, D. Vincent, R. Dadashi, S. Girgin, M. Geist, O. Bachem, O. Pietquin, and M. Andrychowicz. What matters for adversarial imitation learning? *Advances in Neural Information Processing Systems*, 2021.
- [59] A. Osokin, F. Bach, and S. Lacoste-Julien. On structured prediction theory with calibrated convex surrogate losses. In *Proc. of NIPS*, pages 302–313, 2017.
- [60] G. Peyré, M. Cuturi, et al. Computational optimal transport: With applications to data science. *Foundations and Trends® in Machine Learning*, 11(5-6):355–607, 2019.
- [61] D. A. Pomerleau. Efficient training of artificial neural networks for autonomous navigation. *Neural computation*, 1991.
- [62] R. T. Rockafellar and R. J.-B. Wets. *Variational analysis*, volume 317. Springer Science & Business Media, 2009.
- [63] A. M. Rubinov. *Abstract convexity and global optimization*. Springer Science & Business Media, 2000.
- [64] S. Russell. Learning agents for uncertain environments. In *Conference on Computational learning theory*, 1998.

- [65] F. Santambrogio. Optimal transport for applied mathematicians. *Birkäuser, NY*, 55(58-63):94, 2015.
- [66] I. Singer. *Abstract convex analysis*, volume 25. John Wiley & Sons, 1997.
- [67] Y. Song and D. P. Kingma. How to train your energy-based models. *arXiv preprint arXiv:2101.03288*, 2021.
- [68] I. Steinwart. How to compare different loss functions and their risks. *Constructive Approximation*, 26(2):225–287, 2007.
- [69] C. Sutton, A. McCallum, et al. An introduction to conditional random fields. *Foundations and Trends in Machine Learning*, 4(4):267–373, 2012.
- [70] R. S. Sutton and A. G. Barto. *Reinforcement learning: An introduction*. 2018.
- [71] M. J. Wainwright, M. I. Jordan, et al. Graphical models, exponential families, and variational inference. *Foundations and Trends® in Machine Learning*, 1(1–2):1–305, 2008.
- [72] A. L. Yuille and A. Rangarajan. The concave-convex procedure. *Neural computation*, 15(4):915–936, 2003.
- [73] T. Zhang. Statistical analysis of some multi-category large margin classification methods. *Journal of Machine Learning Research*, 5(Oct):1225–1251, 2004.
- [74] T. Zhang. Statistical behavior and consistency of classification methods based on convex risk minimization. *The Annals of Statistics*, 32(1):56–85, 2004.
- [75] X. Zhou. On the fenchel duality between strong convexity and lipschitz continuous gradient. *arXiv preprint arXiv:1803.06573*, 2018.