

# Parametric Majorization for Data-Driven Energy Minimization Methods

Jonas Geiping      Michael Moeller

Department of Electrical Engineering and Computer Science, University of Siegen

{jonas.geiping, michael.moeller}@uni-siegen.de

## Abstract

*Energy minimization methods are a classical tool in a multitude of computer vision applications. While they are interpretable and well-studied, their regularity assumptions are difficult to design by hand. Deep learning techniques on the other hand are purely data-driven, often provide excellent results, but are very difficult to constrain to predefined physical or safety-critical models. A possible combination between the two approaches is to design a parametric energy and train the free parameters in such a way that minimizers of the energy correspond to desired solution on a set of training examples. Unfortunately, such formulations typically lead to bi-level optimization problems, on which common optimization algorithms are difficult to scale to modern requirements in data processing and efficiency. In this work, we present a new strategy to optimize these bi-level problems. We investigate surrogate single-level problems that majorize the target problems and can be implemented with existing tools, leading to efficient algorithms without collapse of the energy function. This framework of strategies enables new avenues to the training of parameterized energy minimization models from large data.*

## 1. Introduction

Energy minimization methods, also referred to as variational methods, are a classical tool in computer vision [74, 14, 27, 12]. The idea is to define a data-dependent cost function  $E$  that assigns a value to each candidate solution  $x$ . The desired optimal solution is then the target solution with the lowest energy value. This methodology has several advantages, for one, it is characterized by an *explicit model* - namely the energy function to be minimized - and an *implicit inference method* - how we compute the minimizer of this energy is a separate problem. This duality allows a fruitful analysis, leading to controllable methods with provable guarantees that are paramount in many critical applications [71, 69, 88]. Furthermore, explicit knowledge over the model structure allows for explainable and clear modifications when the method is applied in a related task [22].

Conversely, deep learning approaches [53], specifically deep feed-forward neural networks work by very different principles. The methodology of deep learning is characterized by *implicit* models and *explicit* inference. The solution to the problem at hand is given directly by the output of the learned feed-forward structure. This is advantageous in practice and crucial for the efficient training of neural networks, however the underlying model of the problem structure is now only implicitly contained in the responses of the network. Deep neural networks have fundamentally changed the state-of-the-art in various computer vision applications, due to these properties as the inference operations are learned directly from large amounts of training data. These approaches are able to learn expressive and convincing mechanisms, examples of which can be found not only in recognition tasks (e.g. [49]), but also in denoising [89], optical flow [63, 43] or segmentation tasks [57, 72, 17]. Yet, as the underlying model is only implicitly defined and 'hidden' in the network structure, it is difficult to modify it for applications in other domains or to guarantee specific outputs. Domain adaptation is still an active field of research and several examples, for instance in medical imaging [3, 33], have demonstrated the need for possibly model-based physically plausible output restrictions. This problem is most strikingly demonstrated by the phenomenon of adversarial examples [79] - the existence of input data, that, when fed through the network, leads to highly erroneous solutions. While one would expect that such behaviour is possibly unavoidable in recognition tasks [77, 64], it should not be a factor in low-level computer vision applications.

Reviewing these two methodologies, we would - of course - prefer to have the best of both worlds. We would like to use both the large amounts of data at our disposal and our far-reaching domain knowledge in many tasks to train explicit models with a significant number of free parameters, so that their optimal solutions are similar to directly trained feed-forward networks.

A promising candidate for such a combination of learning- and model based approaches are *parametrized energy minimization methods*. The idea of such methods is to

define an energy  $E$  that depends on the candidate solutions  $x$ , the input data  $y$  and parameters  $\theta$ ,

$$E : \mathbb{R}^n \times \mathbb{R}^m \times \mathbb{R}^s \rightarrow \mathbb{R}, \quad (1)$$

$$(x, y, \theta) \mapsto E(x, y, \theta),$$

such that for a good choice of parameters  $\theta$ , the argument  $x(\theta) = \arg \min_x E(x, y, \theta)$  that minimizes the energy over all  $x$  is as close as possible to the desired true solution  $x^*$ .

To train such parametric energies, assume we are given  $N$  training samples  $\{(x_i^*, y_i)\}_{i=1}^N$  and a continuous *higher-level* loss function  $l : \mathbb{R}^n \times \mathbb{R}^n \rightarrow \mathbb{R}$ , which measures the deviation of solutions of the model to the given training samples. Determining the optimal parameters  $\theta$  then becomes a *bi-level optimization problem* combining both the higher-level loss function and the lower-level energy,

$$\min_{\theta \in \mathbb{R}^s} \sum_{i=1}^N l(x_i^*, x_i(\theta)), \quad (2)$$

$$\text{subject to } x_i(\theta) = \arg \min_{x \in \mathbb{R}^n} E(x, y_i, \theta). \quad (3)$$

Usual first-order learning methods are difficult to apply in this setting. For every gradient computation it is necessary to compute a derivative of the  $\arg \min$  operation of the lower-level problem, which is even further complicated if we consider parametrized non-smooth energy models which are wide-spread in computer vision [27, 12].

Therefore, the goal of this paper is to analyze bi-level optimization problems and identify strategies that allow for efficient approximate solutions. We investigate single-level minimization problems with simple constraints without second-order differentiation, which are applicable even to non-smooth energies. Such forms allow scaling the previously limited training of energy minimization methods in computer vision to larger datasets and increase the effectiveness in applications where it is critical that the solution follows a specific model structure.

In the remainder of this paper we analyze the bi-level optimization problem to develop a rigorous understanding of sufficient conditions for a single-level surrogate strategy for continuous loss functions  $l$  and convex, non-smooth lower-level energies  $E$  to be successful. We introduce the concept of a *parametric majorization function*, show relations to structured support vector machines and provide several levels of parametric majorization functions with varying levels of exactness and computational effort. We extend our approximations to an iterative scheme, allowing for repeated evaluations of the approximation, before illustrating the proposed strategies in computer vision applications.

## 2. Related Work

The straightforward way of optimizing bi-level problems is to consider *direct descent methods* [48, 76, 26]. These

methods directly differentiate the higher-level loss function with respect to the minimizing argument and descend in the direction of this gradient. An incomplete list of examples in image processing is [11, 22, 20, 21, 28, 29, 36, 40, 41]. This strategy requires both the higher- and lower-level problems to be smooth and the minimizing map to be invertible. This is usually facilitated by implicit differentiation, as discussed in [75, 50, 21, 22]. In more generality, the problem of directly minimizing  $\theta$  without assuming that smoothness in  $E$  leads to optimization problems with equilibrium constraints (MPECs), see [7] for a discussion in terms of machine learning or [31, 30, 32] and [26]. This approach also applies to the optimization layers of [2], which lend themselves well to a reformulation as a bi-level optimization problem.

*Unrolling* is a prominent strategy in applied bi-level optimization across fields, i.e. MRF literature [4, 62] in deep learning [90, 18, 15, 56] and in variational settings [66, 52, 51, 38, 39, 68]. The problem is transformed into a single level problem by choosing an optimization algorithm  $\mathcal{A}$  that produces an approximate solution to the lower level problem after a fixed number of iterations.  $x(\theta)$  is then replaced by  $\mathcal{A}(y, \theta)$ . Automatic differentiation [37] allows for an efficient evaluation of the gradient of the upper-level loss w.r.t to this reduced objective

$$\min_{\theta} \sum_{i=1}^N l(x_i^*, \mathcal{A}(y_i, \theta)). \quad (4)$$

In general these strategies are very successful in practice, *because* they combine the model and its optimization method into a single feed-forward process, where the model is again only implicitly present. Later works [23, 19, 38, 39] allow the lower-level parameters to change in between the fixed number of iterations, leading to structures that model differential equations and stray further from underlying modelling. As pointed out in [46], these strategies are more aptly considered as a set of nested quadratic lower-level problems.

Several techniques have been developed in the field of structured support vector machines (SSVMs) [82, 24, 1, 85] that are very relevant to the task of learning energy models, as SSVMs can be understood as bi-level problems with a lower-level energy that is linear in  $\theta$  and often a non-continuous higher-level loss. Various strategies such as margin rescaling [82], slack rescaling [85, 87], softmax-margins [35] exist and have also been applied recently in the training of computer vision models in [47, 25], we will later return to their connection to the investigated strategies.

## 3. Bi-Level Learning

We now formalize our learning problem. We assume the lower-level energy  $E$  from (1) to be convex (but not necessarily smooth) in its first variable  $x \in \mathbb{R}^n$  and to depend

continuously on input data  $y \in \mathbb{R}^m$  and parameters  $\theta \in \mathbb{R}^s$ . We assume its minimizer  $x(\theta)$  to be unique. For our higher-level loss function (2)  $l : \mathbb{R}^n \times \mathbb{R}^n \rightarrow \mathbb{R}$ , we assume that it fulfills  $l(x, y) \geq 0$ ,  $l(x, x) = 0$  for all  $x, y$  and is differentiable in its second argument.

Note that this formulation of bi-level optimization problems directly generalizes classical supervised (deep) learning with a network  $\mathcal{N}(\theta, y)$  via the quadratic energy  $E(x, y_i, \theta) = \frac{1}{2} \|x - \mathcal{N}(\theta, y_i)\|^2$ , for which  $x_i(\theta) = \mathcal{N}(\theta, y_i)$ .

*Preliminaries (Convex Analysis):* Let us summarize our notation and some fundamental results from convex analysis. We refer the reader to [5] for more details. We denote by  $\partial E(x)$  the set of subgradients of a convex function  $E$  at  $x$ . We define the Bregman distance between two vectors relative to a convex function  $E$  by  $D_E^p(x, y) = E(x) - E(y) - \langle p, x - y \rangle$  for a subgradient  $p \in \partial E(y)$ , intuitively the Bregman distance measures the difference of the energy at  $x$  to its linear lower bound around  $y$ .  $E^*(p) = \sup_x \langle p, x \rangle - E(x)$  is the convex conjugate of  $E$ .  $x$  is a minimizer of the energy  $E$  if and only if  $0 \in \partial E(x)$  or equivalently by convex duality  $x \in \partial E^*(0)$ .  $E$  is  $m$ -strongly convex if  $D_E^p(x, y) \geq \frac{m}{2} \|x - y\|^2$  for all  $x, y$ . Conversely, if  $E$  is  $m$ -strongly convex, then  $E^*$  is  $\frac{1}{m}$ -strongly smooth, i.e.  $D_{E^*}^q(p, q) \leq \frac{2}{m} \|p - q\|^2$ . Furthermore  $D_E^p(x, y) = D_{E^*}^q(p, q)$ ,  $q \in \partial E(x)$  holds for all Bregman distances [9]. We consider parametrized energies in several variables, yet we always assume (sub)-gradients, Bregman distances and convex conjugates to be with respect to the first argument  $x$ .

### 3.1. Majorization of Bi-level Problems

As previously discussed, directly solving the bi-level problem as posed in Eq. (2) and (3) is tricky. We need to implicitly differentiate the minimizing argument  $x_i(\theta)$  for all  $N$  samples just to apply a first-order method in  $\theta$  - which is in stark contrast to our goal of finding efficient and scalable algorithms.

Let us instead look at the problem from a very different angle and entertain the idea that the loss function  $l$  is actually of secondary importance to us. We really only want to find parameters  $\theta$  so that our training samples are well reconstructed,  $x_i^* \approx x_i(\theta)$ . If we go so far as to assume that the loss value of our optimal parameters  $\theta^*$  is zero, meaning that minimizers of our energy are perfectly able to reconstruct our training samples, then the bi-level problem is reduced to a single-level problem, inserting  $x_i^* = x_i(\theta^*)$ :

$$\min_{\theta} \quad \text{s.t. } 0 \in \partial E(x_i^*, y_i, \theta), \quad (5)$$

which we could solve via

$$\min_{\theta} \sum_{i=1}^N \|q_i\|^2 \quad \text{s.t. } q_i \in \partial E(x_i^*, y_i, \theta) \quad (6)$$

This train of thought is closely interconnected to the notion of separability in Support Vector Machine methods [86], where it is assumed that given training samples are linearly separable, which is equivalent to assuming that the classification loss is zero on the training set.

However minimizing Eq. (6) is often not a good choice. A simple example is  $E(x, y, \theta) = (\theta x - y)^2$ , i.e. we simply try to learn a positive scaling factor  $\theta$  between  $x$  and  $y$ . Problem (5) can then be written as  $\min_{\theta} \sum_i (\theta^2 x_i^* - \theta y_i)^2$  and is trivially minimized by  $\theta = 0$ . Such a solution makes  $E$  independent of  $x$  such that every  $x$  becomes a minimizer. This phenomenon is referred to as *collapse* of the energy function [55, 54] in machine learning literature, and clearly cannot be a good strategy to learn a scaling factor.

Interestingly, the scaling problem can be reformulated into a reasonable (non-collapsing) problem, if we require (6) to *majorize* the bilevel problem: If we consider the higher-level loss function  $l(x_i^*, x_i(\theta)) = (x_i^* - x_i(\theta))^2$ , then our surrogate problem  $\sum_i (\theta^2 x_i^* - \theta y_i)^2$  is clearly not a majorizer for arbitrary  $\theta$ . However, if we consider a reformulation of the energy to  $E(x) = (x - \frac{1}{\theta} y)^2$ , then this reformulation leads to a *majorizing* surrogate  $\sum_i (x_i^* - \frac{1}{\theta} y_i)^2$ . Minimizing  $\theta$  now leads to learning the desired scaling factor.

Our toy example motivates us to formalize the concept of majorizing surrogates:

**Definition 1** (Parametrized Majorizer). Given a bi-level optimization problem in the higher level loss  $l(x, y)$  and lower-level energy  $E(x, y, \theta)$ , we call the function  $S(x, y, \theta) : \mathbb{R}^n \times \mathbb{R}^m \times \mathbb{R}^s \rightarrow \mathbb{R}$  a parametrized majorizer, if

$$\begin{aligned} \forall \theta \in \mathbb{R}^s : \quad & l(x, x(\theta)) \leq S(x, y, \theta) \\ \forall \theta \in \mathbb{R}^s \quad \text{s.t.} \quad & l(x, x(\theta)) = 0 \implies S(x, y, \theta) = 0 \end{aligned}$$

hold for any  $x, y \in \mathbb{R}^n \times \mathbb{R}^m$ .

This definition allows us to formalize our objective further. We investigate replacing the bi-level optimization problem (2), (3) by the minimization of a suitable parametrized majorizer, i.e.

$$\min_{\theta \in \mathbb{R}^s} \sum_{i=1}^N S(x_i^*, y_i, \theta). \quad (7)$$

An immediate conclusion of Definition 1 is that the function  $S$  now certifies our progress as  $S(x, y, \theta) = 0$  implies  $l(x, x(\theta)) = 0$ . Moreover, our goal is to choose majorizers  $S$  in such a way that they yield **single-level** problems (7), meaning it is not necessary to differentiate an arg min operation to minimize them or to solve an equally difficult reformulation, making them significantly easier to solve.

### 3.2. Single-Level Majorizers

One possible way to find a majorizer that satisfies the previously postulated properties is by considering the majorizer naturally induced through the Bregman distance of the lower level energy. We assume the following condition

$$l(x, z) \leq D_{E_\theta}(x, z) \quad \forall x, z \in \mathbb{R}^n, \theta \in \mathbb{R}^s, \quad (8)$$

and propose the surrogate problem

$$\min_{\theta} \sum_{i=1}^N D_{E_\theta}(x_i^*, x_i(\theta)). \quad (9)$$

Condition (8) is an assumption on both the loss function and the energy. It thus delineates the class of bi-level problems that can be attacked with this majorization strategy. However this condition is quite general. For a large class of loss functions, we only need the energy to contain a term that also induces the loss function, a property also known as (relative) strong convexity [84, 58]:

**Proposition 1.** *If the loss function  $l(x, y)$  is a Bregman distance induced by a strictly convex function  $w : \mathbb{R}^n \rightarrow \mathbb{R}$ , i.e.  $l(x, y) = D_w(x, y)$ , then assumption (8) is fulfilled if the energy  $E$  is  $w$ -strongly convex, i.e. if  $E(x) - w(x)$  is still a convex function.*

*Proof:* We write  $E$  as  $E(x) = \hat{E}(x) + w(x)$  and apply the additive separability of Bregman distances to find  $D_E(x, y) = D_{\hat{E}}(x, y) + D_w(x, y)$ , which is greater than or equal to  $D_w(x, y)$ , as  $D_{\hat{E}}(x, y)$  is non-negative due to the convexity of  $\hat{E}$ . For the usual euclidean loss, this property reduces to strong convexity:

**Example 1.** *If the loss function is given by a squared Euclidean loss,  $l(x, y) = \frac{1}{2}\|x - y\|^2$  and the energy is  $m$ -strongly convex, then assumption (8) is fulfilled for the energy  $\frac{1}{m}E$ .*

The question remains whether the proposed surrogate problem (9) is efficiently solvable. We especially wanted to circumvent the differentiation of  $x(\theta)$ . However  $D_E(x_i^*, x_i(\theta))$  is much easier to solve, in comparison to the original bi-level problem, as we can see in both its primal and its dual formulation. First, from a primal viewpoint, we have

$$D_E(x_i^*, x_i(\theta)) = E(x_i^*, y_i, \theta) - E(x_i(\theta), y_i, \theta) - \langle p_i, x_i^* - x_i(\theta) \rangle,$$

for some subgradient  $p_i \in \partial E(x_i(\theta))$  which we have not specified yet. But, as  $0 \in \partial E(x_i(\theta))$  as  $x_i(\theta)$  is by definition a solution to the lower-level problem, we may take  $p = 0$  and simplify to

$$E(x_i^*, y_i, \theta) - E(x_i(\theta), y_i, \theta).$$

Now  $x_i(\theta)$  is contained solely in  $E$  and we can write

#### Bregman Surrogate:

$$D_{E_\theta}^0(x_i^*, x_i(\theta)) = \max_{x \in \mathbb{R}^n} E(x_i^*, y_i, \theta) - E(x, y_i, \theta). \quad (10)$$

This surrogate function is already much simpler than the original bi-level problem. We can minimize (10) either by alternating minimization in  $\theta$  and maximization in  $x$  or by jointly optimizing both variables. However, the problem is still set up as a saddle-point problem which is not ideal for optimization.

*Remark.* Interestingly, this discriminative formulation is not wholly unfamiliar. We can understand this as an appropriate generalization of generalized perceptron training [55, 54, 80] as discussed as far back as [73]. See the appendix for further details. In vein of this comparison, conditions 1 and 2 from e.g. [55], i.e. conditions on the existence of a margin between the optimal solution and other candidate solutions central to (S)SVM methods [86, 81, 83] are reflected in Proposition 1 in the convex continuous setting. Due to continuity of the energy and loss function we cannot obey a fixed margin, yet we impose that the energy grows at least as fast as the loss function, when we move away from the optimal solution.

We can resolve the saddle-point question by analyzing the surrogate (9) from a dual standpoint, as by Bregman duality [8]

$$D_{E_\theta}^0(x_i^*, x_i(\theta)) = D_{E_\theta}^{x_i^*}(0, q_i) \quad (11)$$

for  $q_i \in \partial E(x_i^*, y_i, \theta)$ . Contrasting this formulation with our initial goal of penalizing the subgradient as in Eq. (6), we see that the Bregman distance induced by  $E^*$  is the natural 'distance' by which to penalize the subgradient in the sense that penalizing the subgradient at  $x_i^*$  with this generalized distance recovers a majorizing surrogate.

We can further simplify the dual formulation by applying Fenchel's theorem:

$$D_{E_\theta}^{x_i^*}(0, q_i) = E(x_i^*, y_i, \theta) + E^*(0, y_i, \theta). \quad (12)$$

Computing  $E^*(0)$  is exactly as difficult as minimizing  $E$  (as  $E^*(0) = \min_x E(x)$ ), so we need to rewrite this surrogate in a tractable manner. To do so, we assume that  $E$  can be additively decomposed into two parts,

$$E(x, y, \theta) = E_1(x, y, \theta) + E_2(x, y, \theta), \quad (13)$$

where both  $E_1$  and  $E_2$  are convex in their first argument and their convex conjugates are simple to compute. Exploiting that  $E^*(0) = \min_z E_1^*(-z) + E_2^*(z)$  yields

$$D_{E_\theta}^{x_i^*}(0, q_i) = \min_{z \in \mathbb{R}^n} E(x_i^*, y_i, \theta) + E_1^*(-z, y, \theta) + E_2^*(z, y, \theta). \quad (14)$$

In comparison to the primal formulation in Eq (10), we have now reformulated the problem from a saddle point problem (minimizing in  $\theta$  and maximizing in  $x$ ) to a pure minimization problem, which is easier to handle. This is a generalization of the dual formulation discussed in the linear context of SSVMs for example in [81, 83].

However for both variants we still need to handle an auxiliary variable. We can trade some of this computational effort for a weaker majorizer by making specific choices for  $z$  in Eq. (14). To illuminate these choices we introduce the function  $W_E(p, x) = E^*(p) + E(x) - \langle p, x \rangle$  [67, 10], which allows us to write

$$D_{E_\theta^*}^{x_i^*}(0, q_i) = \min_{z \in \mathbb{R}^n} W_{E_1, \theta}(-z, x_i^*) + W_{E_2, \theta}(z, x_i^*). \quad (15)$$

Note that  $W_E(p, x) = 0$  if  $p \in \partial E(x)$ . As such choosing either  $-z \in \partial E_1(x_i^*)$  or  $z \in \partial E_2(x_i^*)$  allows us to simplify the problem further. This is especially attractive if  $E$  is differentiable, as then both surrogates can be computed without auxiliary variables. We will denote these as *partial* surrogates, owing to the fact that we minimize only one term in (15)

#### Partial Surrogate:

$$\min_{z \in \partial E_2(x_i^*, y_i, \theta)} W_{E_1, \theta}(-z, x_i^*). \quad (16)$$

Effectively, this reduces the requirements of (14), as only the convex conjugate of  $E_1$  needs to be computed. By symmetry, the other partial surrogate follows analogously.

We can finally also return to the previously discussed gradient penalty (6). If our energy  $E$  is  $m(\theta, y)$ -strongly convex, then its convex conjugate is strongly smooth and we can bound the dual formulation (11) via

#### Gradient Penalty

$$\frac{1}{m(\theta, y)} \|q_i\|^2 \quad \text{s.t. } q_i \in \partial E(x_i^*, y_i, \theta). \quad (17)$$

While this formulation allows us to minimize an upper bound on the bi-level problem without either auxiliary variables or knowledge about  $E_1^*$  or  $E_2^*$ , it also is the crudest over-approximation among the considered surrogates as the following proposition illustrates.

**Proposition 2** (Ordering of parametric majorizers). *Assuming the condition  $l(x, z) \leq D_{E_\theta}(x, z)$  from Eq. (8), we find that the presented parametric majorizers can be ordered in the following way:*

$$\begin{aligned} l(x_i^*, x(\theta)) &\leq D_{E_\theta}^0(x_i^*, x_i(\theta)) = D_{E_\theta^*}^{x_i^*}(0, q_i) \\ &\leq \min_{z \in \partial E_2(x_i^*)} W_{E_1}(-z, x_i^*) \\ &\leq \frac{1}{m(\theta, y)} \|q_i\|^2 \quad \text{s.t. } q_i \in \partial E(x_i^*, y, \theta). \end{aligned}$$

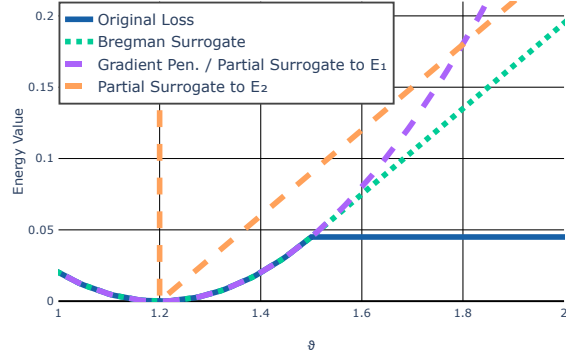


Figure 1. Visualization of surrogate functions for the bi-level problem given in Eq. (18). The blue line marks the original bi-level problem, the green dots marks the Bregman distance surrogate discussed in Eq. (10). The orange curve marks the partial surrogate obtained from (15) by inserting  $z = \nabla E_1(x^*)$ , whereas the purple line marks the other partial surrogate (16) which is equivalent to the gradient penalty (17) here.

The Bregman surrogate (10) majorizes the original loss function and is in turn majorized by the partial surrogate (16) which is majorized by the gradient penalty (17) under the assumption of strong convexity.

*Proof.* See appendix.  $\square$

As a clarifying example, we can simplify these majorizers in the differentiable setting:

**Example 2** (Differentiable Energy). *Let  $E$  be differentiable and  $m(\theta, y)$ -strongly convex, then the majorizers in Prop. 2 are given by*

$$\begin{aligned} l(x_i^*, x(\theta)) &\leq D_{E_\theta}(x_i^*, x_i(\theta)) = D_{E_\theta^*}(0, \nabla E(x_i^*, y_i, \theta)) \\ &\leq W_{E_1}(-\nabla E_2(x_i^*), x_i^*) \\ &\leq \frac{1}{m(\theta, y)} \|\nabla E(x_i^*, y_i, \theta)\|^2. \end{aligned}$$

### 3.3. Intermission: One-Dimensional Example

Let us illustrate our discussion with a toy example. We consider the non-smooth bi-level problem of learning the optimal sparsity parameter  $\theta$  in the bi-level problem:

$$\min_{\theta \in \mathbb{R}} \frac{1}{2} |x^* - x(\theta)|^2, \quad (18)$$

$$\text{subject to } x(\theta) = \arg \min_x \frac{1}{2} |x - y|^2 + \theta |x|. \quad (19)$$

As the lower-level energy is 1-strongly convex and the upper level loss is quadratic  $l(x, y) \leq D_{E_\theta}(x, y)$  holds. Detailed derivations of all three surrogate functions of this example can be found in the appendix. Figure 1 visualizes

these surrogates, plotting their energy values relative to  $\theta$ . Due to the low dimensionality of the problem, all surrogate functions coincide with the original loss function at the optimal value of  $\theta$ . It is further interesting to note that the Bregman surrogate is exactly identical with the original loss function in the vicinity of the optimal value, due to the low dimensionality of the example.

### 3.4. Iterative Majorizers

We used subsection 3.2 to construct a series of upper bounds to facilitate a trade-off between efficiency and exactness. However what happens if we are not satisfied with the exactness of the Bregman surrogate (9)? This setting can happen especially if  $x^*$  and  $x(\theta)$  are significantly incompatible and subsequently  $l(x^*, x(\theta))$  is large, even for optimal  $\theta$ . For example if we try to optimize only a few hyper-parameters we might not at all expect  $x(\theta)$  to be close to  $x^*$ . This discussion can again be linked to the notion of 'separability' in SVM approaches [86]: The quality of the majorizing strategy is directly related to the level of 'separability' of the bi-level problem.

However, we can use the previously introduced majorizers iteratively. To do so we need to develop a majorizer that depends on a given estimate  $\bar{x}$ .

**Proposition 3.** *Under the standing assumption that  $l(x, y) \leq D_{E_\theta}(x, y)$  (8) and if the loss function is induced by a strictly convex function  $w : \mathbb{R}^n \rightarrow \mathbb{R}$ , i.e.  $l(x, y) = D_w(y, x)$ , we have the following inequality:*

$$l(x, y) \leq l(x, z) + \langle \nabla_z l(x, z), y - z \rangle + D_E(z, y). \quad (20)$$

*Proof.* It holds that  $l(x, y) = D_w(y, x)$  which is equivalent to  $D_w(y, z) + D_w(z, x) - \langle \nabla w(x) - \nabla w(z), z - y \rangle$  by the Bregman 3-Point inequality [16, 84]. Using the standing assumption and that  $\nabla w(x) - \nabla w(z) = \nabla_x D_w(x, z)$  we find the proposed inequality.  $\square$

Assume we are given an estimated solution  $\bar{x}_i$ , then we can use this estimate to rewrite our bound to

$$l(x_i^*, x_i(\theta)) \leq l(x_i^*, \bar{x}_i) + \langle \nabla l(x_i^*, \bar{x}_i), x_i(\theta) - \bar{x}_i \rangle + D_E(\bar{x}_i, x_i(\theta)). \quad (21)$$

This is a linearized variant of the parametric majorization bound and as such a nonconvex composite majorizer in the sense of [34], as such a key property of majorization-minimization techniques remains in the parametrized setting, choosing  $\bar{x}_i = x_i(\theta^k)$ :

**Proposition 4** (Descent Lemma). *The iterative procedure given by repeatedly minimizing the right-hand side of Eq. (21) in  $\theta$  and setting  $\bar{x}_i = x_i(\theta^k)$  is guaranteed to be stable, i.e. not to increase the bi-level loss:*

$$\sum_{i=1}^N l(x_i^*, x_i(\theta^{k+1})) \leq \sum_{i=1}^N l(x_i^*, x_i(\theta^k)) \quad (23)$$

*Proof.* See appendix.  $\square$

However this algorithm cannot be applied directly, as we would still need to differentiate  $x_i(\theta)$  appearing in the linearized part. Nevertheless, we can use both Fenchel's inequality  $\langle p, x \rangle \leq E(x) + E^*(p)$  and the previously established  $D_{E_\theta}(x, x(\theta)) = E(x, y, \theta) - E(x(\theta), y, \theta)$  to find an over-approximation to the iterative majorizer of Prop. 4:

$$\begin{aligned} & l(x_i^*, x_i(\theta)) \\ & \leq l(x_i^*, \bar{x}_i) - \langle \nabla l(x_i^*, \bar{x}_i), \bar{x}_i \rangle \\ & \quad + E^*(\nabla l(x_i^*, \bar{x}_i), y_i, \theta) + E(x_i(\theta), y_i, \theta) \\ & \quad + E(\bar{x}_i, y_i, \theta) - E(x_i(\theta), y_i, \theta) \\ & = l(x_i^*, \bar{x}_i) - \langle \nabla l(x_i^*, \bar{x}_i), \bar{x}_i \rangle \\ & \quad + E(\bar{x}_i, y_i, \theta) + E^*(\nabla l(x_i^*, \bar{x}_i), y_i, \theta) \end{aligned}$$

This estimate reveals that we can approximate the iterative majorizer much like the previously discussed surrogates:

#### Iterative Surrogate

$$E(\bar{x}_i, y, \theta) + E^*(\nabla l(x_i^*, \bar{x}_i), y_i, \theta) + C, \quad (22)$$

as the constant  $C = l(x_i^*, \bar{x}_i) - \langle \nabla l(x_i^*, \bar{x}_i), \bar{x}_i \rangle$  does not depend on  $\theta$ . We essentially return to Eq. (12) and only the input to  $E$  and  $E^*$  changes with respect to  $\bar{x}_i$ . This strategy recovers the previous majorizer as a special case:

**Corollary 1.** *If we linearize around  $\bar{x}_i = x_i^*$ , then we recover the Bregman surrogate of (9).*

*Proof.* If  $\bar{x}_i = x_i^*$ , then  $l(x_i^*, \bar{x}_i) = 0$  and  $\nabla l(x_i^*, \bar{x}_i) = 0$  by the properties of the differentiable loss function. As such the constant term  $C$  is zero and  $E^*(\nabla l(x_i^*, \bar{x}_i), y_i, \theta) = E^*(0, y_i, \theta)$  so that we recover (12) which is equivalent to the Bregman surrogate (9).  $\square$

We can use this surrogate to form an efficient approximation to a classical majorization-minimization strategy as in [78, 60, 59, 42]. Notably the 'tightness' of the majorization is violated by the over-approximation, i.e. inserting  $\theta^k$  into the majorizer does not recover  $l(x_i^*, x_i(\theta^k))$ . We iterate

$$\begin{aligned} \theta^{k+1} = \arg \min_{\theta} & \sum_{i=1}^N E^*(\nabla l(x_i^*, x_i(\theta^k)), y_i, \theta) \\ & + E(x(\theta^k), y_i, \theta) \end{aligned} \quad (23)$$

As the application of this iterative scheme reduces to a simple change from Eq (12) to Eq. (22), we can easily apply it in practice to further increase the fidelity of the surrogate by solving a sequence of fast surrogate optimizations. We initialize the scheme with  $\bar{x}_i = x_i^*$  as suggested from Corollary 1 and either stop iterating or reduce the step size of the surrogate solver if the higher-level objective is increased after an iteration.

## 4. Examples

This section will feature several experiments<sup>1</sup> in which we will illustrate the application of the investigated methods. We will show two concepts of new applications that are possible in parametrized variational settings, 4.1 and 4.2. We then show an application to image denoising in 4.3.

### 4.1. Computed Tomography

Making only specific parts of a variational model learnable is especially interesting for computed tomography (CT). An image  $x$  is to be reconstructed from data  $y = Ax + n$  that is formed by applying the radon transform to the image  $x$  and adding noise  $n$ . While first fully-learning based solutions to this problem exist (e.g. [44, 45]), suitable networks are difficult to find not only due to the ill-posedness of the underlying problem, but also due to the well-justified concerns about fully learning-based approaches in medical imaging [3]. To benefit from the explicit control of the data fidelity of the reconstruction, we consider to introduce a learnable linear correction term into an otherwise classical reconstruction technique via

$$x_i(\theta) = \arg \min_x \frac{1}{2} \|Ax - y_i\|_2^2 + \beta R(x) + \langle x, \mathcal{N}(\theta, y_i) \rangle,$$


for a suitable network  $\mathcal{N}$  (we chose 8 blocks of  $3 \times 3$  convolutions with 32 filters, ReLU activations, and batch-normalization, and a final  $5 \times 5$  convolution), and  $R$  denoting the Huber loss of the discrete gradient of  $x$ .

As both convex conjugates are difficult to evaluate in closed-form, we choose the gradient penalty (17), which is a parametric majorizer for euclidean loss if  $A$  has full rank (and practically even works beyond this setting, as it majorizes  $\|A(x - y)\|^2$  even for rank-deficient  $A$ ). According to (17) we consider

$$\min_{\theta \in \mathbb{R}^s} \sum_{i=1}^n \|A^* A x_i^* - A^* y_i + \beta \nabla R(x_i^*) + \mathcal{N}(\theta, y_i)\|_2^2,$$

train on simulated noisy data and test our model on the widely-used Shepp-Logan phantom. Figure 2 illustrates the resulting reconstruction, as well as the best reconstruction using the variational approach without the additional linear correction term after a grid-search for the optimal  $\beta$ . As we can see, the surrogate trained the linear correction term well enough to improve the PSNR of the reconstruction by almost 2dB. Moreover, the influence of the linear correction term can still be visualized and the data fidelity can easily be controlled via a suitable weighting. We visualize the correction map in the appendix.

<sup>1</sup>An implementation of these experiments can be found at <https://github.com/JonasGeiping/ParametricMajorization>.



Huber-TV, PSNR 23.9

Learned cor., PSNR 25.8

Figure 2. Learning a linear correction term for a Huber-regularized CT reconstruction problem using the gradient penalty (6).

### 4.2. Variational Segmentation

For a very different (and non-smooth) example, consider the task of learning a variational segmentation model [14, 13, 27, 65]. We are interested in learning a model whose minimizer coincides with a (semantic) segmentation of the input data. The lower-level problem is given by

$$x(\theta) = \arg \min_x -\langle \mathcal{N}(\theta, y), x \rangle + \|Dx\|_1 + h(x), \quad (24)$$

where  $h(x) = \sum_{j=1}^n x_j \log(x_j) + I_{\Delta}(x)$  is the entropy function on the unit simplex  $\Delta$  [6].  $\mathcal{N}(\theta, y)$  is some parametrized function that computes the potential of the segmentation model, this can be a deep neural network, as we only require convexity in  $x$  and not in  $\theta$ .  $D$  is a finite-differences operator, so that the overall total variation (TV) term  $\|Dx\|_1$  measures the perimeter of a segmentation  $x$  if  $x \in \{0, 1\}^n$ . The entropy function crucially not only leads to a strictly convex model but also represents the structure of a usual learned segmentation method. Without the perimeter term, a solution to the lower-level problem would be given by

$$x(\theta) = \nabla h^*(\mathcal{N}(\theta, y)). \quad (25)$$

Due to [70, P.148],  $\nabla h^*$  is exactly the softmax function, so that Eq. (25) is equivalent to applying a parametrized function  $\mathcal{N}$  and then applying the softmax function to arrive at the final output, a usual image recognition pipeline during training. As a higher-level loss, we choose log loss

$$\sum_{i=1}^N -\langle x_i^*, \log(x_i(\theta)) \rangle = \sum_{i=1}^N D_h(x_i^*, x_i(\theta)) \quad (26)$$

so that the bi-level problem without the perimeter term is equivalent to minimizing the cross-entropy loss of  $\mathcal{N}(\theta, y)$ . With the inclusion of the perimeter term, however, we cannot find a closed-form solution for  $x(\theta)$  need to consider bi-level optimization. But, as the log-loss (26) can be written as a Bregman distance relative to  $h$ , our primary assumption  $l(x, z) \leq D_{E_\theta}(x, z)$  (8) is fulfilled and we can consider the




Figure 3. Training accuracy for the variational segmentation model discussed in Section 4.2 for a linear model  $\mathcal{N}(\theta, y_i)$ . Directly training a cross-entropy loss without the perimeter term, training the Bregman surrogate Eq (28), the Partial surrogate Eq (16) and four iterations of the iterative scheme are compared. We find that the end-to-end training with the perimeter term increases the segmentation accuracy. We also see that a small number of iterations in the iterative scheme is sufficient for a practical CV task.

Bregman surrogate problem in the dual setting of Eq. (14):

$$\min_{\theta} \sum_{i=1}^N \min_{z_i} W_h(\mathcal{N}(\theta, y_i) - z_i, x_i^*) + W_{TV}(z_i, x_i^*), \quad (27)$$

which we can rewrite to

$$\min_{\theta} \sum_{i=1}^N \min_{\|p_i\| \leq 1} h^*(\mathcal{N}(\theta, y_i) - D^T p_i) - \langle \mathcal{N}(\theta, y_i), x_i^* \rangle + \|D x_i^*\|_1. \quad (28)$$

We note that this is essentially a cross-entropy loss with an additional additive term  $p_i$ , that is able to balance out incoherent output of  $\mathcal{N}(\theta, y_i)$  that would lead to erroneous segmentations with a higher perimeter. Furthermore, the training process is still convex w.r.t to  $\mathcal{N}(\theta, y_i)$ , in contrast to unrolling schemes. The iterative model (23) has a very similar structure, including the gradient of the loss into (28).

To validate this setup, we choose  $\mathcal{N}$  to be given by a simple convolutional linear model. We draw a small subset of the `cityscapes` dataset and compare the cross entropy model of Eq (25) with the total variation bi-level model of Eq. (28) and its partial and iterative applications. Figure 3 visualizes the training accuracy over training iterations. We find that the proposed approach is able to improve the segmentation accuracy of the linear model significantly. We refer to the appendix for further details.

### 4.3. Analysis Operator Models

Finally, we illustrate the behaviour of our approach on a practically relevant model, learning a set of optimal convolutional filters for denoising [74, 22]. We consider the

Model	PSNR	T	PSNR(Iter.)	TT
Total Variation	27.41	-	-	-
3 3x3 Filters	26.66	00:34	27.66	02:21
48 7x7 Filters	27.41	02:45	28.03	03:11
96 9x9 Filters	27.46	01:43	28.03	02:22

Table 1. Training time (T) in minutes for each surrogate computation and PSNR on the test dataset for various gray-scale filters for the energy model in Eq. (30) with and without the iterative process of Eq (22) and total time (TT) for the iterative process are compared to total variation with optimal regularization parameter. Note that training time varies mostly due to differing iteration counts. The results of the convex model of [22] are reproduced.

parametric energy model

$$x(\theta) = \arg \min_x \frac{1}{2} \|x - y_i\|^2 + \|D(\theta)x\|_1, \quad (29)$$

with  $D(\theta)$  denoting the convolution operator to be learned, which is prototypical for many other image processing tasks. We consider square loss  $l(x, y) = \frac{1}{2} \|x - y\|^2$  as a higher loss function and apply our approach. A Bregman surrogate for this model has the form

$$\min_{\theta} \sum_{i=1}^N \min_{\|p_i\| \leq 1} \|D(\theta)x_i^*\|_1 + \frac{1}{2} \|D^T(\theta)p_i - y_i\|^2. \quad (30)$$

Model (29) was previously considered in [22, 20], where it was solved via implicit differentiation. We repeat the setup of [22] and train a denoising model on the `BSDS` dataset [61]. Refer to the appendix for the experimental setup and optimization strategy.

Table 1 shows both PSNR values achieved when training  $D(\theta)$  as convolutional filters as well as training time. In comparison to [22], we find strikingly, that we can train a convex model with similar performance to the convex model in [22], while being an order of magnitude faster than the original approach. Furthermore in [22], the necessary training time jumps from 24 hours for 48 7x7 filters to 20 days for 96 9x9 filters - in our experiment the training time is almost unaffected by the number of parameters, and in this example actually smaller as the larger model converges faster. Also this analysis validates that the iterative process is crucial to reaching competitive PSNR values.

## 5. Conclusions

We investigated approximate training strategies for data-driven energy minimization methods by introducing *parametric majorizers*. We systematically studied such strategies in the framework of convex analysis, and proposed the Bregman distance induced by the lower level energy as well as over-approximations thereof as suitable majorizers. We discussed an iterative scheme that shows promise for applications in computer vision, particularly due to its scalability as shown by its application to image denoising.



## References

- [1] Yasemin Altun, Ioannis Tsochantaridis, and Thomas Hofmann. Hidden Markov Support Vector Machines. In *Proceedings of the Twentieth International Conference on International Conference on Machine Learning, ICML'03*, pages 3–10. AAAI Press, 2003. 2
- [2] Brandon Amos and J. Zico Kolter. OptNet: Differentiable Optimization as a Layer in Neural Networks. In *International Conference on Machine Learning*, pages 136–145, July 2017. 2
- [3] Vegard Antun, Francesco Renna, Clarice Poon, Ben Adcock, and Anders C. Hansen. On instabilities of deep learning in image reconstruction - Does AI come at a cost? *arXiv:1902.05300 [cs]*, Feb. 2019. 1, 7
- [4] Adrian Barbu. Learning real-time MRF inference for image denoising. *2009 IEEE Conference on Computer Vision and Pattern Recognition*, pages 1574–1581, 2009. 2
- [5] Heinz H. Bauschke and Patrick L. Combettes. *Convex Analysis and Monotone Operator Theory in Hilbert Spaces*. CMS Books in Mathematics. Springer New York, New York, NY, 2011. 3
- [6] Amir Beck and Marc Teboulle. Mirror descent and nonlinear projected subgradient methods for convex optimization. *Operations Research Letters*, 31(3):167–175, May 2003. 7
- [7] Kristin P. Bennett, Gautam Kunapuli, Jing Hu, and Jong-Shi Pang. Bilevel Optimization and Machine Learning. In *IEEE World Congress on Computational Intelligence, WCCI 2008*, Lecture Notes in Computer Science, pages 25–47. Springer Berlin Heidelberg, Berlin, Heidelberg, 2008. 2
- [8] Martin Benning and Martin Burger. Modern regularization methods for inverse problems. *Acta Numerica*, 27:1–111, May 2018. 4
- [9] Martin Burger. Bregman Distances in Inverse Problems and Partial Differential Equations. In *Advances in Mathematical Modeling, Optimization and Optimal Control*, Springer Optimization and Its Applications, pages 3–33. Springer International Publishing, Cham, 2016. 3
- [10] Dan Butnariu and Gabor Kassay. A Proximal-Projection Method for Finding Zeros of Set-Valued Operators. *SIAM J. Control Optim.*, 47(4):2096–2136, Jan. 2008. 5
- [11] Luca Calatroni, Chung Cao, Juan Carlos De Los Reyes, Carola-Bibiane Schönlieb, and Tuomo Valkonen. Bilevel approaches for learning of variational imaging models. *Variational Methods: In Imaging and Geometric Control*, 18:252, 2017. 2
- [12] Antonin Chambolle, Vicent Caselles, Daniel Cremers, Matteo Novaga, and Thomas Pock. An introduction to total variation for image analysis. *Theoretical foundations and numerical methods for sparse recovery*, 9(263-340):227, 2010. 1, 2
- [13] Antonin Chambolle, Daniel Cremers, and Thomas Pock. A Convex Approach to Minimal Partitions. *SIAM Journal on Imaging Sciences*, 5(4):1113–1158, Oct. 2012. 7
- [14] Tony F. Chan and Luminita A. Vese. Active contours without edges. *IEEE Transactions on image processing*, 10(2):266–277, 2001. 1, 7
- [15] Siddhartha Chandra and Iasonas Kokkinos. Fast, Exact and Multi-scale Inference for Semantic Image Segmentation with Deep Gaussian CRFs. In *Computer Vision – ECCV 2016*, Lecture Notes in Computer Science, pages 402–418. Springer International Publishing, 2016. 2
- [16] Gong Chen and Marc Teboulle. Convergence Analysis of a Proximal-Like Minimization Algorithm Using Bregman Functions. *SIAM J. Optim.*, 3(3):538–543, Aug. 1993. 6
- [17] Liang-Chieh Chen, George Papandreou, Iasonas Kokkinos, Kevin Murphy, and Alan L. Yuille. Semantic Image Segmentation with Deep Convolutional Nets and Fully Connected CRFs. In *International Conference on Learning Representations (ICLR)*, 2015. 1
- [18] Liang-Chieh Chen, George Papandreou, Iasonas Kokkinos, Kevin Murphy, and Alan L. Yuille. DeepLab: Semantic Image Segmentation with Deep Convolutional Nets, Atrous Convolution, and Fully Connected CRFs. *arXiv:1606.00915 [cs]*, June 2016. 2
- [19] Yunjin Chen and Thomas Pock. Trainable Nonlinear Reaction Diffusion: A Flexible Framework for Fast and Effective Image Restoration. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 39(6):1256–1272, June 2017. 2
- [20] Yunjin Chen, Thomas Pock, and Horst Bischof. Learning l1-based analysis and synthesis sparsity priors using bi-level optimization. In *Neural Information Processing Systems Conference (NIPS) 2012*, 2012. 2, 8
- [21] Yunjin Chen, Rene Ranftl, and Thomas Pock. A bi-level view of inpainting - based image compression. In *Computer Vision Winter Workshop.*, 2014. 2
- [22] Yunjin Chen, René Ranftl, and Thomas Pock. Insights Into Analysis Operator Learning: From Patch-Based Sparse Models to Higher Order MRFs. *IEEE Transactions on Image Processing*, 23(3):1060–1072, Mar. 2014. 1, 2, 8
- [23] Yunjin Chen, Wei Yu, and Thomas Pock. On Learning Optimized Reaction Diffusion Processes for Effective Image Restoration. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 5261–5269, 2015. 2
- [24] Michael Collins. Discriminative Training Methods for Hidden Markov Models: Theory and Experiments with Perceptron Algorithms. In *Proceedings of the ACL-02 Conference on Empirical Methods in Natural Language Processing - Volume 10*, EMNLP '02, pages 1–8, Stroudsburg, PA, USA, 2002. Association for Computational Linguistics. 2
- [25] Aleksander Colovic, Patrick Knöbelreiter, Alexander Shekhovtsov, and Thomas Pock. End-to-End Training of Hybrid CNN-CRF Models for Semantic Segmentation using Structured Learning. In *Computer Vision Winter Workshop*, Feb. 2017. 2
- [26] Benoît Colson, Patrice Marcotte, and Gilles Savard. An overview of bilevel optimization. *Annals of Operations Research*, 153(1):235–256, June 2007. 2
- [27] Daniel Cremers, Thomas Pock, Kalin Kolev, and A. Chambolle. Convex Relaxation Techniques for Segmentation, Stereo and Multiview Reconstruction. In *Markov Random Fields for Vision and Image Processing*. MIT Press, Boston, 2011. 1, 2, 7

- [28] Juan Carlos De Los Reyes, Carola-Bibiane Schönlieb, and Tuomo Valkonen. The structure of optimal parameters for image restoration problems. *Journal of Mathematical Analysis and Applications*, 434(1):464–500, Feb. 2016. 2
- [29] Juan Carlos De Los Reyes, Carola-Bibiane Schönlieb, and Tuomo Valkonen. Bilevel Parameter Learning for Higher-Order Total Variation Regularisation Models. *J Math Imaging Vis*, 57(1):1–25, Jan. 2017. 2
- [30] Stephan Dempe. *Foundations of Bilevel Programming*. Non-convex Optimization and Its Applications. Springer US, 2002. 2
- [31] Stephan Dempe and Joydeep Dutta. Is bilevel programming a special case of a mathematical program with complementarity constraints? *Math. Program.*, 131(1):37–48, Feb. 2012. 2
- [32] Stephan Dempe, Vyacheslav Kalashnikov, Gerardo A. Pérez-Valdés, and Nataliya Kalashnykova. *Bilevel Programming Problems: Theory, Algorithms and Applications to Energy Networks*. Energy Systems. Springer-Verlag, Berlin Heidelberg, 2015. 2
- [33] Samuel G. Finlayson, Hyung Won Chung, Isaac S. Kohane, and Andrew L. Beam. Adversarial Attacks Against Medical Deep Learning Systems. *arXiv:1804.05296 [cs, stat]*, Apr. 2018. 1
- [34] Jonas Geiping and Michael Moeller. Composite Optimization by Nonconvex Majorization-Minimization. *SIAM J. Imaging Sci.*, pages 2494–2528, Jan. 2018. 6
- [35] Kevin Gimpel and Noah A. Smith. Softmax-Margin CRFs: Training Log-Linear Models with Cost Functions. In *Human Language Technologies: The 2010 Annual Conference of the North American Chapter of the Association for Computational Linguistics*, pages 733–736, Los Angeles, California, June 2010. Association for Computational Linguistics. 2
- [36] Stephen Gould, Basura Fernando, Anoop Cherian, Peter Anderson, Rodrigo Santa Cruz, and Edison Guo. On Differentiating Parameterized Argmin and Argmax Problems with Application to Bi-level Optimization. *arXiv:1607.05447 [cs, math]*, July 2016. 2
- [37] Andreas Griewank. *Evaluating Derivatives: Principles and Techniques of Algorithmic Differentiation*. Society for Industrial and Applied Mathematics, Philadelphia, PA, USA, 2000. 2
- [38] Kerstin Hammernik, Teresa Klatzer, Erich Kobler, Michael P. Recht, Daniel K. Sodickson, Thomas Pock, and Florian Knoll. Learning a variational network for reconstruction of accelerated MRI data. *Magn Reson Med*, 79(6):3055–3071, June 2018. 2
- [39] Kerstin Hammernik, Tobias Würfl, Thomas Pock, and Andreas Maier. A Deep Learning Architecture for Limited-Angle Computed Tomography Reconstruction. In *Bildverarbeitung für die Medizin 2017*, Informatik aktuell, pages 92–97. Springer Berlin Heidelberg, 2017. 2
- [40] Michael Hintermüller and Carlos N. Rautenberg. Optimal Selection of the Regularization Function in a Weighted Total Variation Model. Part I: Modelling and Theory. *J Math Imaging Vis*, 59(3):498–514, Nov. 2017. 2
- [41] Michael Hintermüller, Carlos N. Rautenberg, Tao Wu, and Andreas Langer. Optimal Selection of the Regularization Function in a Weighted Total Variation Model. Part II: Algorithm, Its Analysis and Numerical Tests. *J Math Imaging Vis*, 59(3):515–533, Nov. 2017. 2
- [42] David R Hunter and Kenneth Lange. A Tutorial on MM Algorithms. *The American Statistician*, 58(1):30–37, Feb. 2004. 6
- [43] Eddy Ilg, Nikolaus Mayer, Tonmoy Saikia, Margret Keuper, Alexey Dosovitskiy, and Thomas Brox. FlowNet 2.0: Evolution of Optical Flow Estimation with Deep Networks. In *2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 1647–1655, July 2017. 1
- [44] Kyong Hwan Jin, Michael T. McCann, Emmanuel Froustey, and Michael Unser. Deep Convolutional Neural Network for Inverse Problems in Imaging. *IEEE Transactions on Image Processing*, 26(9):4509–4522, Sept. 2017. 7
- [45] Eunhee Kang, Junhong Min, and Jong Chul Ye. A deep convolutional neural network using directional wavelets for low-dose X-ray CT reconstruction. *Medical Physics*, 44(10):e360–e375, Oct. 2017. 7
- [46] Teresa Klatzer, Kerstin Hammernik, Patrick Knobelreiter, and Thomas Pock. Learning joint demosaicing and denoising based on sequential energy minimization. In *2016 IEEE International Conference on Computational Photography (ICCP)*, pages 1–11, May 2016. 2
- [47] Patrick Knobelreiter, Christian Reinbacher, Alexander Shekhovtsov, and Thomas Pock. End-to-End Training of Hybrid CNN-CRF Models for Stereo. In *2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 1456–1465, Honolulu, HI, July 2017. IEEE. 2
- [48] Charles. D. Kolstad and Leon S. Lasdon. Derivative evaluation and computational experience with large bilevel mathematical programs. *J Optim Theory Appl*, 65(3):485–499, June 1990. 2
- [49] Alex Krizhevsky, Ilya Sutskever, and Geoffrey E. Hinton. Imagenet classification with deep convolutional neural networks. In *Advances in Neural Information Processing Systems*, pages 1097–1105, 2012. 1
- [50] Karl Kunisch and Thomas Pock. A Bilevel Optimization Approach for Parameter Learning in Variational Models. *SIAM Journal on Imaging Sciences*, 6(2):938–983, Jan. 2013. 2
- [51] Måns Larsson, Anurag Arnab, Fredrik Kahl, Shuai Zheng, and Philip Torr. A Projected Gradient Descent Method for CRF Inference Allowing End-to-End Training of Arbitrary Pairwise Potentials. In *Energy Minimization Methods in Computer Vision and Pattern Recognition*, Lecture Notes in Computer Science, pages 564–579. Springer International Publishing, 2018. 2
- [52] Måns Larsson, Anurag Arnab, Shuai Zheng, Philip Torr, and Fredrik Kahl. Revisiting Deep Structured Models for Pixel-Level Labeling with Gradient-Based Inference. *SIAM J. Imaging Sci.*, pages 2610–2628, Jan. 2018. 2
- [53] Yann LeCun, Yoshua Bengio, and Geoffrey Hinton. Deep learning. *Nature*, 521(7553):436–444, May 2015. 1
- [54] Yann LeCun, Sumit Chopra, Raia Hadsell, M. Ranzato, and F. Huang. A tutorial on energy-based learning. *Predicting structured data*, 1(0), 2006. 3, 4
- [55] Yann LeCun and Fu Jie Huang. Loss functions for discriminative training of energy-based models. In *AISTATS 2005* -

- Proceedings of the 10th International Workshop on Artificial Intelligence and Statistics*, pages 206–213, 2005. 3, 4
- [56] Guosheng Lin, Chunhua Shen, Anton van den Hengel, and Ian Reid. Efficient Piecewise Training of Deep Structured Models for Semantic Segmentation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 3194–3203, 2016. 2
- [57] Jonathan Long, Evan Shelhamer, and Trevor Darrell. Fully Convolutional Networks for Semantic Segmentation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 3431–3440, 2015. 1
- [58] Haihao Lu, Robert M. Freund, and Yurii Nesterov. Relatively Smooth Convex Optimization by First-Order Methods, and Applications. *SIAM J. Optim.*, pages 333–354, Jan. 2018. 4
- [59] Julien Mairal. Optimization with First-order Surrogate Functions. In *Proceedings of the 30th International Conference on International Conference on Machine Learning - Volume 28, ICML'13*, pages III–783–III–791, Atlanta, GA, USA, 2013. JMLR.org. 6
- [60] Julien Mairal. Incremental Majorization-Minimization Optimization with Application to Large-Scale Machine Learning. *SIAM Journal on Optimization*, 25(2):829–855, Jan. 2015. 6
- [61] David Martin, Charless Fowlkes, Doron Tal, and Jitendra Malik. A Database of Human Segmented Natural Images and its Application to Evaluating Segmentation Algorithms and Measuring Ecological Statistics. In *Proceedings of 8th International Conference on Computer Vision*, volume 2, pages 416–423, July 2001. 8
- [62] André F. T. Martins, Noah A. Smith, and Eric P. Xing. Polyhedral outer approximations with application to natural language parsing. In *Proceedings of the 26th Annual International Conference on Machine Learning - ICML '09*, pages 1–8, Montreal, Quebec, Canada, 2009. ACM Press. 2
- [63] Nikolaus Mayer, Eddy Ilg, Philip Häusser, Philipp Fischer, Daniel Cremers, Alexey Dosovitskiy, and Thomas Brox. A Large Dataset to Train Convolutional Networks for Disparity, Optical Flow, and Scene Flow Estimation. *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 4040–4048, June 2016. 1
- [64] Jan Hendrik Metzen, Mummadi Chaithanya Kumar, Thomas Brox, and Volker Fischer. Universal Adversarial Perturbations Against Semantic Image Segmentation. In *2017 IEEE International Conference on Computer Vision (ICCV)*, pages 2774–2783, Oct. 2017. 1
- [65] Claudia Nieuwenhuis, Eno Töppe, and Daniel Cremers. A Survey and Comparison of Discrete and Continuous Multi-label Optimization Approaches for the Potts Model. *Int J Comput Vis*, 104(3):223–240, Sept. 2013. 7
- [66] Peter Ochs, René Ranftl, Thomas Brox, and Thomas Pock. Bilevel Optimization with Nonsmooth Lower Level Problems. In *Scale Space and Variational Methods in Computer Vision*, Lecture Notes in Computer Science, pages 654–665. Springer International Publishing, 2015. 2
- [67] Simeon Reich and Shoham Sabach. Existence and Approximation of Fixed Points of Bregman Firmly Nonexpansive Mappings in Reflexive Banach Spaces. In *Fixed-Point Algorithms for Inverse Problems in Science and Engineering*, Springer Optimization and Its Applications, pages 301–316. Springer, New York, NY, 2011. 5
- [68] Gernot Riegler, Matthias Rütger, and Horst Bischof. Atgynet: Accurate depth super-resolution. In *European Conference on Computer Vision*, pages 268–284. Springer, 2016. 2
- [69] Ludwig Ritschl, Frank Bergner, Christof Fleischmann, and Marc Kachelrie\s s. Improved total variation-based CT image reconstruction applied to clinical data. *Phys. Med. Biol.*, 56(6):1545–1561, Feb. 2011. 1
- [70] R. Tyrrell Rockafellar. *Convex Analysis*. Princeton University Press, Princeton, N.J, 1970. 7
- [71] Torsten Rohlfing, Calvin R. Maurer, David A. Bluemke, and M. A. Jacobs. Volume-preserving nonrigid registration of MR breast images using free-form deformation with an incompressibility constraint. *IEEE Transactions on Medical Imaging*, 22(6):730–741, June 2003. 1
- [72] Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-Net: Convolutional Networks for Biomedical Image Segmentation. In *Medical Image Computing and Computer-Assisted Intervention – MICCAI 2015*, Lecture Notes in Computer Science, pages 234–241. Springer International Publishing, 2015. 1
- [73] Frank Rosenblatt. The perceptron: A probabilistic model for information storage and organization in the brain. *Psychological Review*, 65(6):386–408, 1958. 4
- [74] Leonid I. Rudin, Stanley Osher, and Emad Fatemi. Nonlinear total variation based noise removal algorithms. *Physica D: Nonlinear Phenomena*, 60(1):259–268, Nov. 1992. 1, 8
- [75] Kegan G. G. Samuel and Marshall F. Tappen. Learning optimized MAP estimates in continuously-valued MRF models. In *IEEE Conference on Computer Vision and Pattern Recognition*, Miami, FL, USA, June 2009. IEEE. 2
- [76] Gilles Savard and Jacques Gauvin. The steepest descent direction for the nonlinear bilevel programming problem. *Operations Research Letters*, 15(5):265–272, June 1994. 2
- [77] Ali Shafahi, W. Ronny Huang, Christoph Studer, Soheil Feizi, and Tom Goldstein. Are adversarial examples inevitable? In *ICLR 2019*, New Orleans, Sept. 2018. 1
- [78] Ying Sun, Prabhu Babu, and Daniel P. Palomar. Majorization-Minimization Algorithms in Signal Processing, Communications, and Machine Learning. *IEEE Transactions on Signal Processing*, 65(3):794–816, Feb. 2017. 6
- [79] Christian Szegedy, Wojciech Zaremba, Ilya Sutskever, Joan Bruna, Dumitru Erhan, Ian Goodfellow, and Rob Fergus. Intriguing properties of neural networks. In *arXiv:1312.6199 [Cs]*, Dec. 2013. 1
- [80] Marshall F. Tappen, Ce Liu, Edward H. Adelson, and William T. Freeman. Learning Gaussian Conditional Random Fields for Low-Level Vision. In *2007 IEEE Conference on Computer Vision and Pattern Recognition*, pages 1–8, June 2007. 4
- [81] Ben Taskar, Vassil Chatalbashev, Daphne Koller, and Carlos Guestrin. Learning Structured Prediction Models: A Large Margin Approach. In *Proceedings of the 22nd International Conference on Machine Learning, ICML '05*, pages 896–903, New York, NY, USA, 2005. ACM. 4, 5

- [82] Ben Taskar, Carlos Guestrin, and Daphne Koller. Max-Margin Markov Networks. In *Advances in Neural Information Processing Systems 16*, pages 25–32. MIT Press, 2004. [2](#)
- [83] Ben Taskar, Simon Lacoste-Julien, and Michael I. Jordan. Structured Prediction, Dual Extragradient and Bregman Projections. *Journal of Machine Learning Research*, 7(Jul):1627–1653, 2006. [4](#), [5](#)
- [84] Marc Teboulle. A simplified view of first order methods for optimization. *Math. Program.*, pages 1–30, May 2018. [4](#), [6](#)
- [85] Ioannis Tsochantaridis, Thorsten Joachims, Thomas Hofmann, and Yasemin Altun. Large Margin Methods for Structured and Interdependent Output Variables. *Journal of Machine Learning Research*, 6(Sep):1453–1484, 2005. [2](#)
- [86] Vladimir Vapnik. *Statistical Learning Theory. 1998*, volume 3. Wiley, New York, 1998. [3](#), [4](#), [6](#)
- [87] Sam Wiseman and Alexander M. Rush. Sequence-to-Sequence Learning as Beam-Search Optimization. In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, pages 1296–1306, Austin, Texas, Nov. 2016. Association for Computational Linguistics. [2](#)
- [88] Wei Zhan, Jiachen Li, Yeping Hu, and Masayoshi Tomizuka. Safe and feasible motion generation for autonomous driving via constrained policy net. In *IECON 2017 - 43rd Annual Conference of the IEEE Industrial Electronics Society*, pages 4588–4593, Oct. 2017. [1](#)
- [89] Kai Zhang, Wangmeng Zuo, Yunjin Chen, Deyu Meng, and Lei Zhang. Beyond a Gaussian Denoiser: Residual Learning of Deep CNN for Image Denoising. *IEEE Transactions on Image Processing*, 26(7):3142–3155, July 2017. [1](#)
- [90] Shuai Zheng, Sadeep Jayasumana, Bernardino Romera-Paredes, Vibhav Vineet, Zhizhong Su, Dalong Du, Chang Huang, and Philip H. S. Torr. Conditional Random Fields as Recurrent Neural Networks. *2015 IEEE International Conference on Computer Vision (ICCV)*, pages 1529–1537, Dec. 2015. [2](#)