

Proximal Mean-field for Neural Network Quantization

Thalaiyasingam Ajanthan^{*1}, Puneet K. Dokania², Richard Hartley¹, and Philip H. S. Torr²

¹Australian National University ²University of Oxford

Abstract

Compressing large Neural Networks (NN) by quantizing the parameters, while maintaining the performance is highly desirable due to reduced memory and time complexity. In this work, we cast NN quantization as a discrete labelling problem, and by examining relaxations, we design an efficient iterative optimization procedure that involves stochastic gradient descent followed by a projection. We prove that our simple projected gradient descent approach is, in fact, equivalent to a proximal version of the well-known mean-field method. These findings would allow the decades-old and theoretically grounded research on MRF optimization to be used to design better network quantization schemes. Our experiments on standard classification datasets (MNIST, CIFAR10/100, TinyImageNet) with convolutional and residual architectures show that our algorithm obtains fully-quantized networks with accuracies very close to the floating-point reference networks.

1. Introduction

Despite the success of deep neural networks, they are highly overparametrized, resulting in excessive computational and memory requirements. Compressing such large networks by quantizing the parameters, while maintaining the performance, is highly desirable for real-time applications, or for resource-limited devices.

In Neural Network (NN) quantization, the objective is to learn a network while restricting the parameters to take values from a small discrete set (usually binary) representing quantization levels. This can be formulated as a *discrete labelling problem* where each learnable parameter takes a label from the discrete set and the learning objective is to find the label configuration that minimizes the empirical loss. This is an extremely challenging discrete optimization problem because the number of label configurations grows exponentially with the number of parameters in the network and the loss function is highly non-convex.

Over the past 20 years, similar large-scale discrete la-

labelling problems have been extensively studied under the context of Markov Random Field (MRF) optimization, and many efficient approximate algorithms have been developed [2, 6, 11, 31, 41, 42]. In this work, we take inspiration from the rich literature on MRF optimization, and design an efficient approximate algorithm based on the popular mean-field method [42] for NN quantization.

Specifically, we first formulate NN quantization as a discrete labelling problem. Then, we relax the discrete solution space to a convex polytope and introduce an algorithm to iteratively optimize the first-order Taylor approximation of the loss function over the polytope. This approach is a (stochastic) gradient descent method with an additional projection step at each iteration. For a particular choice of projection, we show that our method is equivalent to a proximal version of the well-known mean-field method. Furthermore, we prove that under certain conditions, our algorithm specializes to the popular BinaryConnect [10] algorithm.

The MRF view of NN quantization opens up many interesting research directions. In fact, our approach represents the simplest case where the NN parameters are assumed to be independent of each other. However, one can potentially model second-order or even high-order interactions among parameters and use efficient inference algorithms developed and well-studied in the MRF optimization literature. Therefore, we believe, many such algorithms can be transposed into this framework to design better network quantization schemes. Furthermore, in contrast to the existing NN quantization methods [21, 35], we quantize *all* the learnable parameters in the network (including biases) and our formulation can be seamlessly extended to quantization levels beyond binary.

We evaluate the merits of our algorithm on MNIST, CIFAR-10/100, and TinyImageNet classification datasets with convolutional and residual architectures. Our experiments show that the quantized networks obtained by our algorithm yield accuracies very close to the floating-point counterparts while consistently outperforming directly comparable baselines. Our code is available at <https://github.com/tajanthan/pmf>.

^{*}Part of the work was done while at the University of Oxford.

2. Neural Network Quantization

Neural Network (NN) quantization is the problem of learning neural network parameters restricted to a small discrete set representing quantization levels. This primarily relies on the hypothesis that overparametrization of NNS makes it possible to obtain a quantized network with performance comparable to the floating-point network. To this end, given a dataset $\mathcal{D} = \{\mathbf{x}_i, \mathbf{y}_i\}_{i=1}^n$, the NN quantization problem can be written as:

$$\min_{\mathbf{w} \in \mathcal{Q}^m} L(\mathbf{w}; \mathcal{D}) := \sum_{i=1}^n \ell(\mathbf{w}; (\mathbf{x}_i, \mathbf{y}_i)). \quad (1)$$

Here, $\ell(\cdot)$ is the input-output mapping composed with a standard loss function (e.g., cross-entropy loss), \mathbf{w} is the m dimensional parameter vector, and \mathcal{Q} with $|\mathcal{Q}| = d$ is a pre-defined discrete set representing quantization levels (e.g., $\mathcal{Q} = \{-1, 1\}$ or $\mathcal{Q} = \{-1, 0, 1\}$). In Eq. (1), we seek a *fully-quantized network* where all the learnable parameters including biases are quantized. This is in contrast to the previous methods [10, 35] where some parts of the network are not quantized (e.g., biases and last layer parameters).

2.1. NN Quantization as Discrete Labelling

NN quantization (1) naturally takes the form of a *discrete labelling problem* where each learnable parameter w_j takes a label q_λ from the discrete set \mathcal{Q} . In particular, Eq. (1) is directly related to an MRF optimization problem [23] where the random variables correspond to the set of weights \mathbf{w} , the label set is \mathcal{Q} , and the energy function is $L(\mathbf{w})$. We refer to Appendix A for a brief overview on MRFs.

An important part of an MRF is the factorization of the energy function that depends on the interactions among the random variables. While modelling a problem as an MRF, emphasis is given to the form of the energy function (e.g., submodularity) as well as the form of the interactions (cliques), because both of these aspects determine the complexity of the resulting optimization. In the case of NNS, the energy function (i.e., loss) is a composition of functions which, in general, has a variety of interactions among the random variables. For example, a parameter at the initial layer is related to parameters at the final layer via function composition. Thus, the energy function does not have an explicit factorization. In fact, optimizing Eq. (1) directly is intractable due to the following inherent problems [26, 32]:

1. The solution space is discrete with exponentially many feasible points (d^m with m in the order of millions).
2. The loss function is highly non-convex and does not satisfy any regularity condition (e.g., submodularity).
3. The loss function does not have an explicit factorization (corresponding to a neighbourhood structure).

This hinders the use of any off-the-shelf discrete optimization algorithm. However, to tackle the aforementioned

problems, we take inspiration from the MRF optimization literature [5, 9, 42]. In particular, we first relax the discrete solution space to a convex polytope and then iteratively optimize the first-order approximation of the loss over the polytope. Our approach, as will be shown subsequently, belongs to the class of (stochastic) gradient descent methods and is applicable to any loss function. Next we describe these relaxations and the related optimization in detail.

2.2. Continuous Relaxation of the Solution Space

Recall that \mathcal{Q} is a finite set of d real-valued parameters. The elements of \mathcal{Q} will be indexed by $\lambda \in \{1, \dots, d\}$. An alternative representation of \mathcal{Q} is by a d -dimensional vector \mathbf{q} with entries $q_\lambda \in \mathcal{Q}$. A element $w \in \mathcal{Q}$ can be written in terms of indicator variables $u_\lambda \in \{0, 1\}$ as $w = \sum_{\lambda=1}^d q_\lambda u_\lambda$, assuming that only one value of q_λ has value 1. Denote by \mathcal{V} the set of size d of such d -vectors with a single 1 component (elements of the standard basis of \mathbb{R}^d) acting as indicator vectors for the elements of \mathcal{Q} . Explicitly, a vector $\mathbf{u}_j \in \mathbb{R}^d$ is in set \mathcal{V} if

$$\sum_{\lambda=1}^d u_{j:\lambda} = 1 \quad \text{and} \quad u_{j:\lambda} \in \{0, 1\} \quad \forall \lambda \in \{1, \dots, d\}.$$

Similarly, the vector $\mathbf{w} \in \mathcal{Q}^m$ of all parameters can be represented using indicator variables as follows. Let $u_{j:\lambda} \in \{0, 1\}$ be the indicator variable, where $u_{j:\lambda} = 1$ if and only if $w_j = q_\lambda \in \mathcal{Q}$. Then, for any $j \in \{1, \dots, m\}$, we can write

$$w_j = \sum_{\lambda=1}^d u_{j:\lambda} q_\lambda = \langle \mathbf{u}_j, \mathbf{q} \rangle \quad \text{where} \quad \mathbf{u}_j \in \mathcal{V}. \quad (2)$$

Any w_j represented using Eq. (2) belongs to \mathcal{Q} . The vector \mathbf{w} of all parameters may be written as a matrix-vector product,

$$\mathbf{w} = \mathbf{u} \mathbf{q} \quad \text{where} \quad \mathbf{u} \in \mathcal{V}^m. \quad (3)$$

Here, $\mathbf{u} = \{u_{j:\lambda}\}$ is thought of as an $m \times d$ matrix (each row \mathbf{u}_j , for $j \in \{1, \dots, m\}$ is an element of \mathcal{V}). Note that there is a one-to-one correspondence between the sets \mathcal{V}^m and \mathcal{Q}^m . Substituting Eq. (3) in the NN quantization objective (1) results in the variable change from \mathbf{w} to \mathbf{u} as:

$$\min_{\mathbf{w} \in \mathcal{Q}^m} L(\mathbf{w}; \mathcal{D}) = \min_{\mathbf{u} \in \mathcal{V}^m} L(\mathbf{u} \mathbf{q}; \mathcal{D}). \quad (4)$$

Even though the above variable change converts the problem from m to md dimensional space, the cardinalities of the sets \mathcal{Q}^m and \mathcal{V}^m are the same. The binary constraint $u_{j:\lambda} \in \{0, 1\}$ together with the non-convex loss function $L(\cdot)$ makes the problem NP-hard [32].

Relaxation. By relaxing the binary constraints to $u_{j:\lambda} \in [0, 1]$, instead of $u_{j:\lambda} \in \{0, 1\}$ we obtain the convex hull Δ^m of the set \mathcal{V}^m . The minimization Eq. (4) may be carried out over Δ^m instead of \mathcal{V}^m . In detail, we define

$$\Delta = \left\{ \mathbf{z} \in \mathbb{R}^d \mid \sum_{\lambda} z_\lambda = 1, z_\lambda \geq 0, \forall \lambda \right\}. \quad (5)$$

This is the standard $(d-1)$ -dimensional simplex embedded in \mathbb{R}^d and the vertices of Δ are the points in \mathcal{V} . Similarly, the Cartesian product Δ^m is the convex hull of \mathcal{V}^m , which are in turn the vertices of Δ^m .

Simplex Δ will be referred to as the *probability simplex* because an element $u \in \Delta$ may be thought of (formally) as a probability distribution on the finite set $\{1, \dots, d\}$. A value u_λ is the probability of choosing the discrete parameter $w = q_\lambda \in \mathcal{Q}$. With this probabilistic interpretation, one verifies that $\mathbf{u}\mathbf{q} = \mathbb{E}_{\mathbf{u}}[\mathbf{w}]$, the expected value of the vector of parameters \mathbf{w} , where each w_j has independent probability distribution defined by \mathbf{u}_j .

Now, the relaxed optimization can be written as:

$$\min_{\mathbf{u} \in \Delta^m} \tilde{L}(\mathbf{u}; \mathcal{D}) := L(\mathbf{u}\mathbf{q}; \mathcal{D}), \quad (6)$$

The minimum of this problem will generally be less than the minimum of Eq. (4). However, if $\mathbf{u} \in \mathcal{V}^m$, then the loss function $\tilde{L}(\mathbf{u})$ has the same value as the original loss function $L(\mathbf{w})$. Furthermore, the relaxation of \mathbf{u} from \mathcal{V}^m to Δ^m translates into relaxing \mathbf{w} from \mathcal{Q}^m to the convex region $[q_{\min}, q_{\max}]^m$. Here, q_{\min} and q_{\max} represent the minimum and maximum quantization levels, respectively.

In fact, $\mathbf{u} \in \Delta^m$ is an overparametrized representation of $\mathbf{w} \in [q_{\min}, q_{\max}]^m$, and the mapping $\mathbf{u} \rightarrow \mathbf{w} = \mathbf{u}\mathbf{q}$ is a many-to-one *surjective* mapping. In the case where $d = 2$ (two quantization levels), the mapping is one-to-one and subjective. In addition it can be shown that any local minimum of Eq. (6) (the relaxed \mathbf{u} -space) is also a local minimum of the loss in $[q_{\min}, q_{\max}]^m$ (the relaxed \mathbf{w} -space) and vice versa (Proposition 2.1). This essentially means that the variable change from \mathbf{w} to \mathbf{u} does not alter the optimization problem and a local minimum in the \mathbf{w} -space can be obtained by optimizing in the \mathbf{u} -space.

Proposition 2.1. Let $f(\mathbf{w}) : [q_{\min}, q_{\max}]^m \rightarrow \mathbb{R}$ be a function, and \mathbf{w} a point in $[q_{\min}, q_{\max}]^m$ such that $\mathbf{w} = g(\mathbf{u}) = \mathbf{u}\mathbf{q}$. Then \mathbf{u} is a local minimum of $f \circ g$ in Δ^m if and only if \mathbf{w} is a local minimum of f in $[q_{\min}, q_{\max}]^m$.

Proof. The function $g : \Delta^m \rightarrow [q_{\min}, q_{\max}]^m$ is surjective continuous and affine. It follows that it is also an open map. From this the result follows easily. \square

Finally, we would like to point out that the relaxation used while moving from \mathbf{w} to \mathbf{u} space is well studied in the MRF optimization literature and has been used to prove bounds on the quality of the solutions [9, 25]. In the case of NN quantization, in addition to the connection to mean-field (Sec. 3), we believe that this relaxation allows for exploration, which would be useful in the stochastic setting.

2.3. First-order Approximation and Optimization

Here we talk about the optimization of $\tilde{L}(\mathbf{u})$ over Δ^m , discuss how our optimization scheme allows exploration in

the parameter space, and also discuss the conditions when this optimization will lead to a quantized solution in the \mathbf{w} space, which is our prime objective.

Stochastic Gradient Descent (SGD)¹ [37] is the de facto method of choice for optimizing neural networks. In this section, we interpret SGD as a proximal method, which will be useful later to show its difference to our final algorithm. In particular, SGD (or gradient descent) can be interpreted as iteratively minimizing the first-order Taylor approximation of the loss function augmented by a proximal term [33]. In our case, the objective function is the same as SGD but the feasible points are now constrained to form a convex polytope. Thus, at each iteration k , the first-order objective can be written as:

$$\begin{aligned} \mathbf{u}^{k+1} &= \operatorname{argmin}_{\mathbf{u} \in \Delta^m} \tilde{L}(\mathbf{u}^k) + \langle \mathbf{g}^k, \mathbf{u} - \mathbf{u}^k \rangle + \frac{1}{2\eta} \|\mathbf{u} - \mathbf{u}^k\|^2, \\ &= \operatorname{argmin}_{\mathbf{u} \in \Delta^m} \langle \mathbf{u}, \eta \mathbf{g}^k - \mathbf{u}^k \rangle + \|\mathbf{u}\|^2/2, \end{aligned} \quad (7)$$

where $\eta > 0$ is the learning rate and $\mathbf{g}^k := \nabla_{\mathbf{u}} \tilde{L}^k$ is the stochastic (or mini-batch) gradient of \tilde{L} with respect to \mathbf{u} evaluated at \mathbf{u}^k . In the unconstrained case, by setting the derivative with respect to \mathbf{u} to zero, one verifies that the above formulation leads to standard SGD updates $\mathbf{u}^{k+1} = \mathbf{u}^k - \eta \mathbf{g}^k$. For constrained optimization (as in our case (7)), it is natural to use the stochastic version of Projected Gradient Descent (PGD) [38]. Specifically, at iteration k , the projected stochastic gradient update can be written as:

$$\mathbf{u}^{k+1} = P_{\Delta^m}(\mathbf{u}^k - \eta \mathbf{g}^k), \quad (8)$$

where $P_{\Delta^m}(\cdot)$ denotes the projection to the polytope Δ^m . Even though this type of problem can be optimized using projection-free algorithms [3, 13, 27], by relying on PGD, we enable the use of any off-the-shelf first-order optimization algorithms (e.g., Adam [24]). Furthermore, for a particular choice of projection, we show that the PGD update is equivalent to a proximal version of the mean-field method.

2.3.1 Projection to the Polytope Δ^m

Projection to Δ^m can be decomposed into m independent projections to the d -dimensional probability simplexes. The objective function (7) is also separable for each j . Thus, for notational convenience, without loss of generality, we assume $m = 1$. Now, for a given updated parameter $\tilde{\mathbf{u}}^{k+1} = \mathbf{u}^k - \eta \mathbf{g}^k$ (where $\tilde{\mathbf{u}}^{k+1} \in \mathbb{R}^d$), we discuss three approaches of projecting to the probability simplex Δ . An illustration of these projections is shown in Fig. 1. In this section, for brevity, we also ignore the superscript $k+1$.

Euclidean Projection (Sparsemax). The standard approach of projecting to a set in the Euclidean space is via sparsemax [30]. Given a scalar $\beta > 0$ (usually $\beta = 1$),

¹The difference between SGD and gradient descent is that the gradients are approximated using a stochastic oracle in the former case.




Figure 1: Illustration of \mathbf{w} and \mathbf{u} -spaces, different projections, and exploration with softmax when $m = 1$. Here each vertex of the simplex corresponds to a discrete quantization level in the \mathbf{w} -space and the simplex is partitioned based on its vertex association. Given an infeasible point $\tilde{\mathbf{u}}$, it is projected to the simplex via softmax (or sparsemax) and when $\beta \rightarrow \infty$, the projected point would move towards the associated vertex.

sparsemax amounts to finding a point \mathbf{u} in Δ which is the closest to $\beta\tilde{\mathbf{u}}$, namely

$$\mathbf{u} = \text{sparsemax}(\beta\tilde{\mathbf{u}}) = \underset{\mathbf{z} \in \Delta}{\operatorname{argmin}} \|\mathbf{z} - \beta\tilde{\mathbf{u}}\|^2. \quad (9)$$

As the name suggests, this projection is likely to hit the boundary of the simplex², resulting in sparse solutions (\mathbf{u}) at every iteration. Please refer to [30] for more detail. As β increases, the projected point moves towards a vertex.

Hardmax Projection. The hardmax projection maps a given $\tilde{\mathbf{u}}$ to one of the vertices of the simplex Δ :

$$\mathbf{u} = \text{hardmax}(\tilde{\mathbf{u}}), \quad (10)$$

$$u_\lambda = \begin{cases} 1 & \text{if } \lambda = \underset{\mu \in \mathcal{Q}}{\operatorname{argmax}} \tilde{u}_\mu \\ 0 & \text{otherwise} \end{cases} \quad \text{for } \lambda \in \{1, \dots, d\}.$$

Softmax Projection. We now discuss the softmax projection which projects a point to the interior of the simplex, leading to dense solutions. Given a scalar $\beta > 0$, the softmax projection is:

$$\mathbf{u} = \text{softmax}(\beta\tilde{\mathbf{u}}), \quad (11)$$

$$u_\lambda = \frac{\exp(\beta\tilde{u}_\lambda)}{\sum_{\mu \in \mathcal{Q}} \exp(\beta\tilde{u}_\mu)} \quad \forall \lambda \in \{1, \dots, d\}.$$

Even though approximate in the Euclidean sense, softmax shares many desirable properties to sparsemax [30] (for example, it preserves the relative order of $\tilde{\mathbf{u}}$) and when $\beta \rightarrow \infty$, the projected point moves towards a vertex.

2.3.2 Exploration and Quantization using Softmax

All of the projections discussed above are valid in the sense that the projected point lies in the simplex Δ . However,

²Unless $\beta\tilde{\mathbf{u}}$ when projected to the simplex plane is in Δ , which is rare.

our goal is to obtain a quantized solution in the \mathbf{w} -space which is equivalent to obtaining a solution \mathbf{u} that is a vertex of the simplex Δ . Below we provide justifications behind using softmax with a monotonically increasing schedule for β in realizing this goal, rather than either sparsemax or hardmax projection.

Recall that the main reason for relaxing the feasible points to lie within the simplex Δ is to simplify the optimization problem with the hope that optimizing this relaxation will lead to a better solution. However, in case of hardmax and sparsemax projections, the effective solution space is restricted to be either the set of vertices \mathcal{V} (no relaxation) or the boundary of the simplex (much smaller subset of Δ). Such restrictions hinder exploration over the simplex and do not fully utilize the potential of the relaxation. In contrast, softmax allows *exploration* over the entire simplex and a monotonically increasing schedule for β ensures that the solution gradually approaches a vertex. This interpretation is illustrated in Fig. 1.

Entropy based view of Softmax. In fact, softmax can be thought of as a “noisy” projection to the set of vertices \mathcal{V} , where the noise is controlled by the hyperparameter β . We now substantiate this interpretation by providing an entropy based view for the softmax projection.

Lemma 2.1. *Let $\mathbf{u} = \text{softmax}(\beta\tilde{\mathbf{u}})$ for some $\tilde{\mathbf{u}} \in \mathbb{R}^d$ and $\beta > 0$. Then,*

$$\mathbf{u} = \underset{\mathbf{z} \in \Delta}{\operatorname{argmax}} \langle \tilde{\mathbf{u}}, \mathbf{z} \rangle + \frac{1}{\beta} H(\mathbf{z}), \quad (12)$$

where $H(\mathbf{z}) = -\sum_{\lambda=1}^d z_\lambda \log z_\lambda$ is the entropy.

Proof. This can be proved by writing the Lagrangian and setting the derivatives to zero. \square

The softmax projection translates into an entropy term in the objective function (12), and for small values of β , it allows the iterative procedure to explore the optimization landscape. We believe, in the stochastic setting, such an explorative behaviour is crucial, especially in the early stage of training. Furthermore, our empirical results validate this hypothesis that PGD with softmax projection is relatively easy to train and yields consistently better results compared to other PGD variants. Note that, when $\beta \rightarrow \infty$, the entropy term vanishes and softmax approaches hardmax.

Note, constraining the solution space through a hyperparameter (β in our case) has been extensively studied in the optimization literature and one such example is the barrier method [7]. Moreover, even though the softmax based PGD update yields an approximate solution to Eq. (7), in Sec. 3, we prove that it is theoretically equivalent to a proximal version of the mean-field method.

3. Softmax based PGD as Proximal Mean-field

Here we discuss the connection between softmax based PGD and the well-known mean-field method [42]. Precisely, we show that the update $\mathbf{u}^{k+1} = \text{softmax}(\beta(\mathbf{u}^k - \eta \mathbf{g}^k))$ is actually an *exact fixed point update* of a modified mean-field objective function. This connection bridges the gap between the MRF optimization and the NN quantization literature. We now begin with a brief review of the mean-field method and then proceed with our proof.

Mean-field Method. A self-contained overview is provided in Appendix A, but here we review the important details. Given an energy (or loss) function $L(\mathbf{w})$ and the corresponding probability distribution of the form $P(\mathbf{w}) = e^{-L(\mathbf{w})}/Z$, mean-field approximates $P(\mathbf{w})$ using a fully-factorized distribution $U(\mathbf{w}) = \prod_{j=1}^m U_j(w_j)$. Here, the distribution U is obtained by minimizing the KL-divergence $\text{KL}(U||P)$. Note that, from the probabilistic interpretation of $\mathbf{u} \in \Delta^m$ (see Sec. 2.2), for each $j \in \{1, \dots, m\}$, the probability $U_j(w_j = q_\lambda) = u_{j:\lambda}$. Therefore, the distribution U can be represented using the variables $\mathbf{u} \in \Delta^m$, and hence, the mean-field objective can be written as:

$$\text{argmin}_{\mathbf{u} \in \Delta^m} \text{KL}(\mathbf{u}||P) = \text{argmin}_{\mathbf{u} \in \Delta^m} \mathbb{E}_{\mathbf{u}}[L(\mathbf{w})] - H(\mathbf{u}), \quad (13)$$

where $\mathbb{E}_{\mathbf{u}}[\cdot]$ is expectation over \mathbf{u} and $H(\mathbf{u})$ is the entropy.

In fact, mean-field has been extensively studied in the MRF literature where the energy function $L(\mathbf{w})$ factorizes over small subsets of variables \mathbf{w} . This leads to efficient minimization of the KL-divergence as the expectation $\mathbb{E}_{\mathbf{u}}[L(\mathbf{w})]$ can be computed efficiently. However, in a standard neural network, the function $L(\mathbf{w})$ does not have an explicit factorization and direct minimization of the KL-divergence is not straight forward. To simplify the NN loss function one can approximate it using its first-order Taylor approximation which discards the interactions between the NN parameters altogether.

In Theorem 3.1, we show that our softmax based PGD iteratively applies a proximal version of mean-field to the first-order approximation of $L(\mathbf{w})$. At iteration k , let $\hat{L}^k(\mathbf{w})$ be the first-order Taylor approximation of $L(\mathbf{w})$. Then, since there are no interactions among parameters in $\hat{L}^k(\mathbf{w})$, and it is linear, our proximal mean-field objective has a closed form solution, which is exactly the softmax based PGD update.

The following theorem applies to the update of each $\mathbf{u}_j \in \Delta$ separately, and hence the update of the corresponding parameter w_j .

Theorem 3.1. *Let $L(\mathbf{u}) : \Delta \rightarrow \mathbb{R}$ be a differentiable function defined in an open neighbourhood of the polytope Δ , and \mathbf{u}^k a point in Δ . Let \mathbf{g}^k be the gradient of $L(\mathbf{u})$ at \mathbf{u}^k , and $\hat{L}^k(\mathbf{u}) = L(\mathbf{u}^k) + \langle \mathbf{u} - \mathbf{u}^k, \mathbf{g}^k \rangle$ the first-order approximation of L at \mathbf{u}^k . Let β and η (learning rate) be positive constants, and*

$$\mathbf{u}^{k+1} = \text{softmax}(\beta(\mathbf{u}^k - \eta \mathbf{g}^k)), \quad (14)$$

the softmax-based PGD update. Then,

$$\mathbf{u}^{k+1} = \text{argmin}_{\mathbf{u} \in \Delta} \eta \hat{L}^k(\mathbf{u}) - \langle \mathbf{u}^k, \mathbf{u} \rangle - \frac{1}{\beta} H(\mathbf{u}). \quad (15)$$

Proof. First one shows that

$$\eta \hat{L}^k(\mathbf{u}) - \langle \mathbf{u}^k, \mathbf{u} \rangle = - \langle \mathbf{u}, \mathbf{u}^k - \eta \mathbf{g}^k \rangle, \quad (16)$$

apart from constant terms (those not containing \mathbf{u}). Then the proof follows from Lemma 2.1. \square

The objective function Eq. (15) is essentially the same as the mean-field objective (13) for $\hat{L}^k(\mathbf{w})$ (noting $\mathbb{E}_{\mathbf{u}}[\hat{L}^k(\mathbf{w})] = \hat{L}^k(\mathbf{u}\mathbf{q}) = \langle \mathbf{g}^k, \mathbf{u} \rangle$ up to constant terms) except for the term $\langle \mathbf{u}^k, \mathbf{u} \rangle$. This, in fact, acts as a proximal term. Note, it is the cosine similarity but subtracted from the loss to enforce proximity. Therefore, it encourages the resulting \mathbf{u}^{k+1} to be closer to the current point \mathbf{u}^k and its influence relative to the loss term is governed by the learning rate η . Since gradient estimates are stochastic in our case, such a proximal term is highly desired as it encourages the updates to make a smooth transition.

Furthermore, the negative entropy term acts as a convex regularizer and when $\beta \rightarrow \infty$ its influence becomes negligible and the update results in a binary labelling $\mathbf{u} \in \mathcal{V}^m$. In addition, the entropy term in Eq. (15) captures the (in)dependency between the parameters. To encode dependency, the entropy of the fully-factorized distribution can perhaps be replaced with a more complex entropy such as a tree-structured entropy, following the idea of [36]. Furthermore, in place of \hat{L}^k , a higher-order approximation can be used. Such explorations go beyond the scope of this paper.

Remark. Note that, our update (15) can be interpreted as an *entropic penalty method* and it is similar in spirit to that of the mirror-descent algorithm when entropy is chosen as the mirror-map (refer Sec. 4.3 of [8]). In fact, at each iteration, both our algorithm and mirror-descent augment the gradient descent objective with a negative entropy term and optimizes over the polytope. However, compared to mirror-descent, our update additionally constitutes a proximal term and an annealing hyperparameter β which enables us to gradually enforce a discrete solution. Therefore, to employ mirror-descent, one needs to understand the effects of using adaptive mirror-maps (that depend on β). Nevertheless, it is interesting to explore the potential of mirror-descent which could allow us to derive different variants of our algorithm.

Proximal Mean-Field (PMF). The preferred embodiment of our PMF algorithm is similar to softmax based PGD. Algorithm 1 summarizes our approach. Similar to the existing methods [21], however, we introduce the auxiliary variables $\tilde{\mathbf{u}} \in \mathbb{R}^{m \times d}$ and perform gradient descent on them, composing the loss function \tilde{L} with the softmax function that maps $\tilde{\mathbf{u}}$ into Δ^m . In effect this solves the optimization problem:

$$\min_{\tilde{\mathbf{u}} \in \mathbb{R}^{m \times d}} \tilde{L}(\text{softmax}(\beta \tilde{\mathbf{u}}); \mathcal{D}). \quad (17)$$

Algorithm 1 Proximal Mean-Field (PMF)

Require: $K, b, \{\eta^k\}, \rho > 1, \mathcal{D}, \tilde{L}$ **Ensure:** $\mathbf{w}^* \in \mathcal{Q}^m$

- 1: $\tilde{\mathbf{u}}^0 \in \mathbb{R}^{m \times d}, \beta \leftarrow 1$ ▷ Initialization
 - 2: **for** $k \leftarrow 0, \dots, K$ **do**
 - 3: $\mathbf{u}^k \leftarrow \text{softmax}(\beta \tilde{\mathbf{u}}^k)$ ▷ Projection (Eq. (11))
 - 4: $\mathcal{D}^b = \{(\mathbf{x}_i, \mathbf{y}_i)\}_{i=1}^b \sim \mathcal{D}$ ▷ Sample a mini-batch
 - 5: $\mathbf{g}_{\mathbf{u}}^k \leftarrow \nabla_{\mathbf{u}} \tilde{L}(\mathbf{u}; \mathcal{D}^b) \Big|_{\mathbf{u}=\mathbf{u}^k}$ ▷ Gradient w.r.t. \mathbf{u} at \mathbf{u}^k
 - 6: $\mathbf{g}_{\tilde{\mathbf{u}}}^k \leftarrow \mathbf{g}_{\mathbf{u}}^k \frac{\partial \mathbf{u}}{\partial \tilde{\mathbf{u}}} \Big|_{\tilde{\mathbf{u}}=\tilde{\mathbf{u}}^k}$ ▷ Gradient w.r.t. $\tilde{\mathbf{u}}$ at \mathbf{u}^k
 - 7: $\tilde{\mathbf{u}}^{k+1} \leftarrow \tilde{\mathbf{u}}^k - \eta^k \mathbf{g}_{\tilde{\mathbf{u}}}^k$ ▷ Gradient descent on $\tilde{\mathbf{u}}$
 - 8: $\beta \leftarrow \rho \beta$ ▷ Increase β
 - 9: **end for**
 - 10: $\mathbf{w}^* \leftarrow \text{hardmax}(\tilde{\mathbf{u}}^K) \mathbf{q}$ ▷ Quantization (Eq. (10))
-

In this way, optimization is carried out over the unconstrained domain $\mathbb{R}^{m \times d}$ rather than over the domain Δ^m . In contrast to existing methods, this is not a necessity but empirically it improves the performance. Finally, since β can never be ∞ , to ensure a fully-quantized network, the final quantization is performed using hardmax. Since, softmax approaches hardmax when $\beta \rightarrow \infty$, the fixed points of Algorithm 1 corresponds to the fixed points of PGD with the hardmax projection. However, exploration due to softmax allows our algorithm to converge to fixed points with better validation errors as demonstrated in the experiments.

3.1. Proximal ICM as a Special Case

For PGD, if hardmax is used instead of the softmax projection, the resulting update is the same as a proximal version of Iterative Conditional Modes (ICM) [5]. In fact, following the proof of Lemma 2.1, it can be shown that the update $\mathbf{u}^{k+1} = \text{hardmax}(\mathbf{u}^k - \eta \mathbf{g}^k)$ yields a fixed point of the following equation:

$$\min_{\mathbf{u} \in \Delta^m} \eta \langle \mathbf{g}^k, \mathbf{u} \rangle - \langle \mathbf{u}^k, \mathbf{u} \rangle. \quad (18)$$

Notice, this is exactly the same as the ICM objective augmented by the proximal term. In this case, $\mathbf{u} \in \mathcal{V}^m \subset \Delta^m$, meaning, the feasible domain is restricted to be the vertices of the polytope Δ^m . Since softmax approaches hardmax when $\beta \rightarrow \infty$, this is a special case of proximal mean-field.

3.2. BinaryConnect as Proximal ICM

In this section, considering binary neural networks, *i.e.*, $\mathcal{Q} = \{-1, 1\}$, and non-stochastic setting, we show that the Proximal Iterative Conditional Modes (PICM) algorithm is equivalent to the popular BinaryConnect (BC) method [10]. In these algorithms, the gradients are computed in two different spaces and therefore to alleviate any discrepancy we assume that gradients are computed using the full dataset.

Let $\tilde{\mathbf{w}} \in \mathbb{R}^m$ and $\mathbf{w} \in \mathcal{Q}^m$ be the infeasible and feasible points of BC. Similarly, $\tilde{\mathbf{u}} \in \mathbb{R}^{m \times d}$ and $\mathbf{u} \in \mathcal{V}^m \subset \Delta^m$

Algorithm 2 One iteration of BinaryConnect (BC) [10]

Require: $\tilde{\mathbf{w}}^k, \eta_{\mathbf{w}}, \mathcal{D}, L$

- 1: $\mathbf{w}^k \leftarrow \text{sign}(\tilde{\mathbf{w}}^k)$ ▷ Projection
 - 2: $\mathbf{g}_{\mathbf{w}}^k \leftarrow \nabla_{\mathbf{w}} L(\mathbf{w}; \mathcal{D}) \Big|_{\mathbf{w}=\mathbf{w}^k}$ ▷ Gradient w.r.t. \mathbf{w}
 - 3: $\mathbf{g}_{\tilde{\mathbf{w}}}^k \leftarrow \mathbf{g}_{\mathbf{w}}^k \frac{\partial \mathbf{w}}{\partial \tilde{\mathbf{w}}} \Big|_{\tilde{\mathbf{w}}=\tilde{\mathbf{w}}^k}$ ▷ Gradient w.r.t. $\tilde{\mathbf{w}}$
 - 4: $\tilde{\mathbf{w}}^{k+1} \leftarrow \tilde{\mathbf{w}}^k - \eta_{\mathbf{w}} \mathbf{g}_{\tilde{\mathbf{w}}}^k$ ▷ Gradient descent
-

be the infeasible and feasible points of our PICM method, respectively. For convenience, we summarize one iteration of BC in Algorithm 2. Now, we show that the update steps in both BC and PICM are equivalent.

Proposition 3.1. Consider BC and PICM with $\mathbf{q} = [-1, 1]^T$ and $\eta_{\mathbf{w}} > 0$. For an iteration $k > 0$, if $\tilde{\mathbf{w}}^k = \tilde{\mathbf{u}}^k \mathbf{q}$ then,

1. the projections in BC: $\mathbf{w}^k = \text{sign}(\tilde{\mathbf{w}}^k)$ and PICM: $\mathbf{u}^k = \text{hardmax}(\tilde{\mathbf{u}}^k)$ satisfy $\mathbf{w}^k = \mathbf{u}^k \mathbf{q}$.
2. let the learning rate of PICM be $\eta_{\mathbf{u}} = \eta_{\mathbf{w}}/2$, then the updated points after the gradient descent step in BC and PICM satisfy $\tilde{\mathbf{w}}^{k+1} = \tilde{\mathbf{u}}^{k+1} \mathbf{q}$.

Proof. Case (1) is simply applying $\tilde{\mathbf{w}}^k = \tilde{\mathbf{u}}^k \tilde{\mathbf{q}}$ whereas case (2) can be proved by writing \mathbf{w}^k as a function of $\tilde{\mathbf{u}}^k$ and then applying chain rule. See Appendix B. \square

Since hardmax is a non-differentiable operation, the partial derivative $\partial \mathbf{u} / \partial \tilde{\mathbf{u}} = \partial \text{hardmax} / \partial \tilde{\mathbf{u}}$ is not defined. However, to allow backpropagation, we write hardmax in terms of the sign function, and used the straight-through-estimator [17] to allow gradient flow similar to binary connect. For details please refer to Appendix B.1.

4. Related Work

There is much work on NN quantization focusing on different aspects such as quantizing parameters [10], activations [20], loss aware quantization [18] and quantization for specialized hardware [12], to name a few. Here we give a brief summary of latest works and for a comprehensive survey we refer the reader to [15].

In this work, we consider parameter quantization, which can either be treated as a post-processing scheme [14] or incorporated into the learning process. Popular methods [10, 21] falls into the latter category and optimize the constrained problem using some form of projected stochastic gradient descent. In contrast to projection, quantization can also be enforced using a penalty term [4, 43]. Even though, our method is based on projected gradient descent, by optimizing in the \mathbf{u} -space, we provide theoretical insights based on mean-field and bridge the gap between NN quantization and MRF optimization literature.

In contrast, the variational approach can also be used for quantization, where the idea is to learn a posterior probability on the network parameters in a Bayesian framework.

Dataset	Image	# class	Train / Val.	b	K
MNIST	28×28	10	50k / 10k	100	20k
CIFAR-10	32×32	10	45k / 5k	128	100k
CIFAR-100	32×32	100	45k / 5k	128	100k
TinyImageNet	64×64	200	100k / 10k	128	100k

Table 1: *Experiment setup. Here, b is the batch size and K is the total number of iterations used for all the methods.*

In this family of methods, the quantized network can be obtained either via a quantizing prior [1] or using the MAP estimate on the learned posterior [40]. Interestingly, the learned posterior distribution can be used to estimate the model uncertainty and in turn determine the required precision for each network parameter [29]. Note that, even in our seemingly different method, we learn a probability distribution over the parameters (see Sec. 2.2) and it would be interesting to understand the connection between Bayesian methods and our algorithm.

5. Experiments

Since neural network binarization is the most popular quantization [10, 35], we set the quantization levels to be binary, *i.e.*, $\mathcal{Q} = \{-1, 1\}$. However, our formulation is applicable to any predefined set of quantization levels given sufficient resources at training time. We would like to point out that, we quantize all learnable parameters, meaning, all quantization algorithms result in 32 times less memory compared to the floating point counterparts.

We evaluate our Proximal Mean-Field (PMF) algorithm on MNIST, CIFAR-10, CIFAR-100 and TinyImageNet³ classification datasets with convolutional and residual architectures and compare against the BC method [10] and the latest algorithm ProxQuant (PQ) [4]. Note that BC and PQ constitute the closest and directly comparable baselines to PMF. Furthermore, many other methods have been developed based on BC by relaxing some of the constraints, *e.g.*, layer-wise scalars [35], and we believe, similar extensions are possible with our method as well. Our results show that the binary networks obtained by PMF yield accuracies very close to the floating point counterparts while consistently outperforming the baselines.

5.1. Experiment Setup

The details of the datasets and their corresponding experiment setups are given in Table 1. In all the experiments, standard multi-class cross-entropy loss is minimized. MNIST is tested using LeNet-300 and LeNet-5, where the former consists of three fully-connected (FC) layers

³<https://tiny-imagenet.herokuapp.com/>

while the latter is composed of two convolutional and two FC layers. For CIFAR and TinyImageNet, VGG-16 [39] and ResNet-18 [16] architectures adapted for CIFAR dataset are used. In particular, for CIFAR experiments, similar to [28], the size of the FC layers of VGG-16 is set to 512 and no dropout layers are employed. For TinyImageNet, the stride of the first convolutional layer of ResNet-18 is set to 2 to handle the image size [19]. In all the models, batch normalization [22] (with no learnable parameters) and ReLU non-linearity are used. Except for MNIST, standard data augmentation is used (*i.e.*, random crop and horizontal flip) and weight decay is set to 0.0001 unless stated otherwise.

For all the algorithms, the hyperparameters such as the optimizer and the learning rate (also its schedule) are cross-validated using the validation set⁴ and the chosen parameters are given in the supplementary material. For PMF and PGD with sparsemax, the growth-rate ρ in Algorithm 1 (the multiplicative factor used to increase β) is cross validated between 1.01 and 1.2 and chosen values for each experiment are given in supplementary. Furthermore, since the original implementation of BC do not binarize all the learnable parameters, for fair comparison, we implemented BC in our experiment setting based on the publicly available code⁵. However, for PQ we used the original code⁶, *i.e.*, for PQ, *biases and last layer parameters are not binarized*. All methods are trained from a random initialization and the model with the best validation accuracy is chosen for each method. Our algorithm is implemented in PyTorch [34].

5.2. Results

The classification accuracies (top-1) on the test set of all versions of our algorithm, namely, PMF, PGD (this is PGD with the sparsemax projection), and PICM, the baselines BC and PQ, and the floating point Reference Network (REF) are reported in Table 2. The training curves for CIFAR-10 and CIFAR-100 with ResNet-18 are shown in Fig. 2. Note that our PMF algorithm consistently produces better results than other binarization methods and the degradation in performance to the full floating point reference network is minimal especially for small datasets. For larger datasets (*e.g.*, CIFAR-100), binarizing ResNet-18 results in much smaller degradation compared to VGG-16.

The superior performance of PMF against BC, PICM and PGD empirically validates the hypothesis that performing “noisy” projection via softmax and annealing the noise is indeed beneficial in the stochastic setting. Furthermore, even though PICM and BC are theoretically equivalent in the non-stochastic setting, PICM yields slightly better accuracies in all our experiments. We conjecture that this is due to

⁴For TinyImageNet, since the ground truth labels for the test set were not available, validation set is used for both cross-validation and testing.

⁵<https://github.com/itayhubara/BinaryNet.pytorch>

⁶<https://github.com/allenbai01/ProxQuant>

Dataset	Architecture	REF (Float)	BC [10]	PQ [4]	Ours			REF - PMF
					PICM	PGD	PMF	
MNIST	LeNet-300	98.55	98.05	98.13	98.18	98.21	98.24	+0.31
	LeNet-5	99.39	99.30	99.27	99.31	99.28	99.44	-0.05
CIFAR-10	VGG-16	93.01	86.40	90.11	88.96	88.48	90.51	+2.50
	ResNet-18	94.64	91.60	92.32	92.02	92.60	92.73	+1.91
CIFAR-100	VGG-16	70.33	43.70	55.10	45.65	57.83	61.52	+8.81
	ResNet-18	73.85	69.93	68.35	70.85	70.60	71.85	+2.00
TinyImageNet	ResNet-18	56.41	49.33	49.97	49.66	49.60	51.00	+5.63

Table 2: Classification accuracies on the test set for different methods. Note that our PMF algorithm consistently produces better results than other binarization methods and the degradation in performance to the full floating point network (last column) is minimal especially for small datasets. For larger datasets (e.g., CIFAR-100), binarizing ResNet-18 results in much smaller degradation compared to VGG-16. Even though, PICM and BC are theoretically equivalent in the non-stochastic setting, PICM yields slightly better accuracies. Note, all binarization methods except PQ require exactly **32** times less memory compared to single-precision floating points networks at test time.




Figure 2: Training curves for CIFAR-10 (first two) and CIFAR-100 (last two) with ResNet-18. For quantization methods, the validation accuracy is always measured with the quantized networks. Specifically, for PMF and PGD, the hardmax projection is applied before the evaluation. Notably, validation accuracy plots clearly illustrate the exploration phase of both PMF and PGD, during which the accuracies are the worst. However, once β is “large enough”, the curves closely resembles high-precision reference network while yielding very high accuracies. Furthermore, compared to BC and PICM, other methods are less noisy suggesting the usefulness of optimizing over a convex domain.

the fact that in PICM, the training is performed on a larger network (*i.e.*, in the \mathbf{u} -space).

To further consolidate our implementation of BC, we quote the accuracies reported in the original papers here. In [10], the top-1 accuracy on CIFAR-10 with a modified VGG type network is 90.10%. In the same setting, even with additional layer-wise scalars, (Binary Weight Network (BWN) [35]), the corresponding accuracy is 90.12%. For comprehensive results on network quantization we refer the reader to Table 5 of [15]. Note that, in all the above cases, the last layer parameters and biases in all layers were not binarized.

6. Discussion

In this work, we have formulated NN quantization as a discrete labelling problem and introduced a projected stochastic gradient descent algorithm to optimize it. By showing our approach as a proximal mean-field method, we have also provided an MRF optimization perspective to

NN quantization. This connection opens up interesting research directions primarily on considering dependency between the neural network parameters to derive better network quantization schemes. Furthermore, our PMF approach learns a probability distribution over the network parameters, which is similar in spirit to Bayesian deep learning methods. Therefore, we believe, it is interesting to explore the connection between Bayesian methods and our algorithm, which can potentially drive research in both fields.

7. Acknowledgements

This work was supported by the ERC grant ERC-2012-AdG 321162-HELIOS, EPSRC grant Seebibyte EP/M013774/1, EPSRC/MURI grant EP/N019474/1 and the Australian Research Council Centre of Excellence for Robotic Vision (project number CE140100016). We would also like to acknowledge the Royal Academy of Engineering, FiveAI, National Computing Infrastructure, Australia and Nvidia (for GPU donation).

References

- [1] J. Achterhold, J. M. Kohler, A. Schmeink, and T. Genewein. Variational network quantization. *ICLR*, 2018. 7
- [2] Thalaiyasingam Ajanthan. *Optimization of Markov random fields in computer vision*. PhD thesis, Australian National University, 2017. 1
- [3] Thalaiyasingam Ajanthan, Alban Desmaison, Rudy Bunel, Mathieu Salzmann, Philip H S Torr, and M Pawan Kumar. Efficient linear programming for dense CRFs. *CVPR*, 2017. 3
- [4] Yu Bai, Yu-Xiang Wang, and Edo Liberty. Proxquant: Quantized neural networks via proximal operators. *ICLR*, 2019. 6, 7, 8
- [5] Julian Besag. On the statistical analysis of dirty pictures. *Journal of the Royal Statistical Society.*, 1986. 2, 6
- [6] Andrew Blake, Pushmeet Kohli, and Carsten Rother. *Markov random fields for vision and image processing*. Mit Press, 2011. 1
- [7] Stephen Boyd and Lieven Vandenberghe. *Convex optimization*. Cambridge university press, 2009. 4
- [8] Sébastien Bubeck. Convex optimization: Algorithms and complexity. *Foundations and Trends® in Machine Learning*, 2015. 5
- [9] Chandra Chekuri, Sanjeev Khanna, Joseph Naor, and Leonid Zosin. A linear programming formulation and approximation algorithms for the metric labeling problem. *SIAM Journal on Discrete Mathematics*, 2004. 2, 3
- [10] Matthieu Courbariaux, Yoshua Bengio, and Jean-Pierre David. Binaryconnect: Training deep neural networks with binary weights during propagations. *NIPS*, 2015. 1, 2, 6, 7, 8
- [11] P. K. Dokania and P. K. Mudigonda. Parsimonious labeling. *ICCV*, 2015. 1
- [12] S. K. Esser, R. Appuswamy, P. A. Merolla, J. V. Arthur, and D. S. Modha. Backpropagation for energy-efficient neuromorphic computing. *NIPS*, 2015. 6
- [13] Marguerite Frank and Philip Wolfe. An algorithm for quadratic programming. *Naval research logistics quarterly*, 1956. 3
- [14] Y. Gong, L. Liu, and L. Bourdev. Compressing deep convolutional networks using vector quantization. *arXiv preprint arXiv:1412.6115*, 2014. 6
- [15] Yunhui Guo. A survey on methods and theories of quantized neural networks. *arXiv preprint arXiv:1808.04752*, 2018. 6, 8
- [16] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. *CVPR*, 2016. 7
- [17] Geoffrey Hinton. Neural networks for machine learning. *Coursera, video lectures*, 2012. 6
- [18] Lu Hou, Quanming Yao, and James T Kwok. Loss-aware binarization of deep networks. *ICLR*, 2017. 6
- [19] Gao Huang, Yixuan Li, Geoff Pleiss, Zhuang Liu, John E Hopcroft, and Kilian Q Weinberger. Snapshot ensembles: Train 1, get m for free. *ICLR*, 2017. 7
- [20] Itay Hubara, Matthieu Courbariaux, Daniel Soudry, Ran El-Yaniv, and Yoshua Bengio. Binarized neural networks. *NIPS*, 2016. 6
- [21] Itay Hubara, Matthieu Courbariaux, Daniel Soudry, Ran El-Yaniv, and Yoshua Bengio. Quantized neural networks: Training neural networks with low precision weights and activations. *JMLR*, 2017. 1, 5, 6
- [22] Sergey Ioffe and Christian Szegedy. Batch normalization: Accelerating deep network training by reducing internal covariate shift. *ICML*, 2015. 7
- [23] R. Kindermann and J. L. Snell. *Markov Random Fields and Their Applications*. American Mathematical Society, 1980. 2
- [24] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *ICLR*, 2015. 3
- [25] Jon Kleinberg and Eva Tardos. Approximation algorithms for classification problems with pairwise relationships: metric labeling and Markov random fields. *Journal of the ACM*, 2002. 3
- [26] Vladimir Kolmogorov and Ramin Zabini. What energy functions can be minimized via graph cuts? *PAMI*, 2004. 2
- [27] Simon Lacoste-Julien, Martin Jaggi, Mark Schmidt, and Patrick Pletscher. Block-coordinate Frank-Wolfe optimization for structural SVMs. *ICML*, 2012. 3
- [28] Namhoon Lee, Thalaiyasingam Ajanthan, and Philip H S Torr. SNIP: Single-shot network pruning based on connection sensitivity. *ICLR*, 2019. 7
- [29] C. Louizos, K. Ullrich, and M. Welling. Bayesian compression for deep learning. *NIPS*, 2017. 7
- [30] Andre Martins and Ramon Astudillo. From softmax to sparsemax: A sparse model of attention and multi-label classification. *ICML*, 2016. 3, 4
- [31] Pawan Kumar Mudigonda. *Combinatorial and convex optimization for probabilistic models in computer vision*. PhD thesis, Oxford Brookes University, 2008. 1
- [32] George L Nemhauser and Laurence A Wolsey. *Integer programming and combinatorial optimization*. Springer, 1988. 2
- [33] Neal Parikh and Stephen P Boyd. Proximal algorithms. *Foundations and Trends in Optimization*, 2014. 3
- [34] Adam Paszke, Sam Gross, Soumith Chintala, Gregory Chanan, Edward Yang, Zachary DeVito, Zeming Lin, Alban Desmaison, Luca Antiga, and Adam Lerer. Automatic differentiation in PyTorch. 2017. 7
- [35] Mohammad Rastegari, Vicente Ordonez, Joseph Redmon, and Ali Farhadi. Xnor-net: Imagenet classification using binary convolutional neural networks. *ECCV*, 2016. 1, 2, 7, 8
- [36] Pradeep Ravikumar, Alekh Agarwal, and Martin J Wainwright. Message-passing for graph-structured linear programs: proximal projections, convergence and rounding schemes. *ICML*, 2008. 5
- [37] Herbert Robbins and Sutton Monro. A stochastic approximation method. *Annals of Mathematical Statistics*, 1951. 3
- [38] Lorenzo Rosasco, Silvia Villa, and Bang Công Vũ. Convergence of stochastic proximal gradient algorithm. *arXiv preprint arXiv:1403.5074*, 2014. 3

- [39] Karen Simonyan and Andrew Zisserman. Very deep convolutional networks for large-scale image recognition. *ICLR*, 2015. [7](#)
- [40] D. Soudry, I. Hubara, and R. Meir. Expectation backpropagation: Parameter-free training of multilayer neural networks with continuous or discrete weights. *NIPS*, 2018. [7](#)
- [41] Olga Veksler. *Efficient graph-based energy minimization methods in computer vision*. PhD thesis, Cornell University New York, USA, 1999. [1](#)
- [42] Martin J Wainwright, Michael I Jordan, et al. Graphical models, exponential families, and variational inference. *Foundations and Trends® in Machine Learning*, 2008. [1](#), [2](#), [5](#)
- [43] Penghang Yin, Shuai Zhang, Jiancheng Lyu, Stanley Osher, Yingyong Qi, and Jack Xin. Binaryrelax: A relaxation approach for training deep neural networks with quantized weights. *SIIMS*, 2018. [6](#)