
Demystifying and Generalizing BinaryConnect

Tim Dockhorn *
University of Waterloo

Yaoliang Yu
University of Waterloo

Eyyüb Sari
Huawei Noah’s Ark Lab

Mahdi Zolnouri
Huawei Noah’s Ark Lab

Vahid Partovi Nia
Huawei Noah’s Ark Lab

Abstract

BinaryConnect (BC) and its many variations have become the de facto standard for neural network quantization. However, our understanding of the inner workings of BC is still quite limited. We attempt to close this gap in four different aspects: (a) we show that existing quantization algorithms, including post-training quantization, are surprisingly similar to each other; (b) we argue for proximal maps as a natural family of quantizers that is both easy to design and analyze; (c) we refine the observation that BC is a special case of dual averaging, which itself is a special case of the generalized conditional gradient algorithm; (d) consequently, we propose *ProxConnect* (PC) as a generalization of BC and we prove its convergence properties by exploiting the established connections. We conduct experiments on CIFAR-10 and ImageNet, and verify that PC achieves competitive performance.

1 Introduction

Scaling up to extremely large datasets and models has been a main ingredient for the success of deep learning. Indeed, with the availability of big data, more computing power, convenient software, and a bag of training tricks as well as algorithmic innovations, the size of models that we routinely train in order to achieve state-of-the-art performance has exploded, e.g., to billions of parameters in recent language models [9]. However, high memory usage and computational cost at inference time has made it difficult to deploy these models in real-time or on resource-limited devices [28]. The environmental impact of training and deploying these large models has also been recognized [40]. A common approach to tackle these problems is to compress a large model through quantization, i.e., replacing high-precision parameters with lower-precision ones. For example, we may constrain a subset of the weights to be binary [1, 5, 11, 29, 33, 37, 45] or ternary [26, 53]. Quantization can drastically decrease the carbon footprint of training and inference of neural networks, however, it may come at the cost of increased bias [21].

One of the main methods to obtain quantized neural networks is to encourage quantized parameters during gradient training using explicit or implicit regularization techniques, however, other methods are possible [16–18, 20, 25, 32, 51]. Besides the memory benefits, the structure of the quantization can speed up inference using, for example, faster matrix-vector products [18, 23]. Training and inference can be made even more efficient by also quantizing the activations [37] or gradients [52]. Impressive performance has been achieved with quantized networks, for example, on object detection [44] and natural language processing [43] tasks. The theoretical underpinnings of quantized neural networks, such as when and why their performance remains reasonably well, have been actively studied [3, 13, 22, 41, 46].

*Work done during an internship at Huawei Noah’s Ark Lab. Correspondence to tim.dockhorn@uwaterloo.ca.

BinaryConnect [BC, 11] and its many variations [10, 37, 53] are considered the gold standard for neural network quantization. Compared to plain (stochastic) gradient descent, BC does not evaluate the gradient at the current iterate but rather at a (close-by) quantization point using the Straight Through Estimator [6]. Despite its empirical success, BC has largely remained a “training trick” [1] and a rigorous understanding of its inner workings has yet to be found, with some preliminary steps taken in Li et al. [27] for the convex setting and in Yin et al. [45] for a particular quantization set. As pointed out in Bai et al. [5], BC only evaluates gradients at the finite set of quantization points, and therefore does not exploit the rich information carried in the continuous weights network. Bai et al. [5] also observed that BC is formally equivalent to the dual averaging algorithm [31, 42], while some similarity to the mirror descent algorithm was found in Ajanthan et al. [1].

The main goal of this work is to significantly improve our understanding of BC, by connecting it with well-established theory and algorithms. In doing so we not only simplify and improve existing results but also obtain novel generalizations. We summarize our main contributions in more details:

- In Section 2, we show that existing gradient-based quantization algorithms are surprisingly similar to each other: the only high-level difference is at what points we evaluate the gradient and perform the update.
- In Section 3, we present a principled theory for constructing proximal quantizers. Our results unify previous efforts, remove tedious calculations, and bring theoretical convenience. We illustrate our theory by effortlessly designing a new quantizer that can be computed in one-pass, works for different quantization sets (binary or multi-bit), and includes previous attempts as special cases [5, 11, 45, 53].
- In Section 4, we significantly extend the observation of Bai et al. [5] that the updates of BC are the same as the dual averaging algorithm [31, 42]: BC is a nonconvex counterpart of dual averaging, and more importantly, dual averaging itself is simply the generalized conditional gradient algorithm applied to a smoothed dual problem. The latter fact, even in the convex case, does not appear to be widely recognized to the best of our knowledge.
- In Section 5, making use of the above established results, we propose *ProxConnect* (PC) as a family of algorithms that generalizes BC and we prove its convergence properties for both the convex and the nonconvex setting. We rigorously justify the diverging parameter in proximal quantizers and resolve a discrepancy between theory and practice in the literature [1, 5, 45].
- In Section 6, we verify that PC outperforms BC and ProxQuant [5] on CIFAR-10 for both fine-tuning pretrained models as well as end-to-end training. On the more challenging ImageNet dataset, PC yields competitive performance despite of minimal hyperparameter tuning.

2 Background

We consider the usual (expected) objective $\ell(\mathbf{w}) = \mathbf{E}\ell(\mathbf{w}, \boldsymbol{\xi})$, where $\boldsymbol{\xi}$ represents random sampling. For instance, $\ell(\mathbf{w}) = \frac{1}{n} \sum_{i=1}^n \ell_i(\mathbf{w})$ where $\boldsymbol{\xi}$ is uniform over n training samples (or minibatches thereof), \mathbf{w} are the weights of a neural network and ℓ_i may be the cross-entropy loss (of the i -th training sample). We denote a sample (sub)gradient of ℓ at \mathbf{w} and $\boldsymbol{\xi}$ as $\tilde{\nabla}\ell(\mathbf{w}) = \nabla\ell(\mathbf{w}, \boldsymbol{\xi})$ so that $\mathbf{E}\tilde{\nabla}\ell(\mathbf{w}) = \nabla\ell(\mathbf{w})$. Throughout, we use the starred notation \mathbf{w}^* for continuous weights and reserve \mathbf{w} for (semi)discrete ones.

We are interested in solving the following (nonconvex) problem:

$$\min_{\mathbf{w} \in Q} \ell(\mathbf{w}), \tag{1}$$

where $Q \subseteq \mathbb{R}^d$ is a discrete, nonconvex quantization set. For instance, on certain low-resource devices it may be useful or even mandatory to employ binary weights, i.e., $Q = \{\pm 1\}^d$. Importantly, our goal is to compete against the non-quantized, continuous weights network (i.e. $Q = \mathbb{R}^d$). In other words, we do *not* necessarily aim to solve (the hard, combinatorial) problem (1) globally and optimally. Instead, we want to find discrete weights $\mathbf{w} \in Q$ that remain satisfactory when compared to the non-quantized continuous weights. This is how we circumvent the difficulty in (1) and more importantly how the structure of ℓ could come into aid. If ℓ is reasonably smooth, a close-by quantized weight of a locally, or even globally, optimal continuous weight will likely yield similar performance.

Tremendous progress has been made on quantizing and compressing neural networks. While it is not possible to discuss all related work, below we recall a few families of gradient-based quantization algorithms that directly motivate our work; more methods can be found in recent surveys [15, 36].

BinaryConnect (BC). Courbariaux et al. [11] considered binary networks where $Q = \{\pm 1\}^d$ and proposed the BinaryConnect algorithm:

$$\mathbf{w}_{t+1}^* = \mathbf{w}_t^* - \eta_t \tilde{\nabla} \ell(\mathbf{P}(\mathbf{w}_t^*)), \quad (2)$$

where \mathbf{P} is a projector that quantizes the continuous weights \mathbf{w}_t^* either deterministically (by taking the sign) or stochastically. Note that the (sample) gradient is evaluated at the quantized weights $\mathbf{w}_t := \mathbf{P}(\mathbf{w}_t^*)$, while its continuous output, after scaled by the step size η_t , is added to the continuous weights \mathbf{w}_t^* . For later comparison, it is useful to break the BC update into the following two pieces:

$$\mathbf{w}_t = \mathbf{P}(\mathbf{w}_t^*), \quad \mathbf{w}_{t+1}^* = \mathbf{w}_t^* - \eta_t \tilde{\nabla} \ell(\mathbf{w}_t). \quad (3)$$

Other choices of \mathbf{P} [1, 45] and Q [44, 53] in this framework have also been experimented with.

ProxQuant (PQ). Bai et al. [5] applied the usual proximal gradient to solve (1):

$$\mathbf{w}_{t+1} = \mathbf{P}(\mathbf{w}_t - \eta_t \tilde{\nabla} \ell(\mathbf{w}_t)), \quad (4)$$

which can be similarly decomposed into:

$$\mathbf{w}_t = \mathbf{P}(\mathbf{w}_t^*), \quad \mathbf{w}_{t+1}^* = \mathbf{w}_t - \eta_t \tilde{\nabla} \ell(\mathbf{w}_t). \quad (5)$$

Thus, the only high-level difference between BC and PQ is that the former updates the continuous weights \mathbf{w}_t^* while the latter updates the quantized weights \mathbf{w}_t . This seemingly minor difference turns out to cause drastically different behaviors of the two algorithms. For example, choosing \mathbf{P} to be the Euclidean projection to Q works well for BinaryConnect but not at all for ProxQuant.

Reversing BinaryConnect. The above comparison naturally suggests a reversed variant of BC:

$$\mathbf{w}_t = \mathbf{P}(\mathbf{w}_t^*), \quad \mathbf{w}_{t+1}^* = \mathbf{w}_t - \eta_t \tilde{\nabla} \ell(\mathbf{w}_t^*), \quad (6)$$

which amounts to switching the continuous and quantized weights in the BC update. Similar to PQ, the choice of \mathbf{P} is critical in this setup. To the best of our knowledge, this variant has not been formally studied before. Reversing BinaryConnect may not be intuitive, however, it serves as a starting point for a more general method (see Section 5).

Post-Training Quantization. Lastly, we can also rewrite the naive post-training quantization scheme in a similar form:

$$\mathbf{w}_t = \mathbf{P}(\mathbf{w}_t^*), \quad \mathbf{w}_{t+1}^* = \mathbf{w}_t^* - \eta_t \tilde{\nabla} \ell(\mathbf{w}_t^*), \quad (7)$$

where we simply train the continuous network as usual and then quantize at the end. Note that the quantized weights \mathbf{w}_t do not affect the update of the continuous weights \mathbf{w}_t^* .

3 What Makes a Good Quantizer?

As we have seen in Section 2, the choice of the quantizer \mathbf{P} turns out to be a crucial element for solving (1). Indeed, if $\mathbf{P} = \mathbf{P}_Q$ is the projector onto the discrete quantization set Q , then BC (and ProxQuant) only evaluate the gradient of ℓ at (the finite set of) points in Q . As a result, the methods will not be able to exploit the rich information carried in the continuous weights network, which can lead to non-convergence [Fig. 1b, 5]. Since then, many semi-discrete quantizers, that turn continuous weights into more and more discrete ones, have been proposed [1, 5, 33, 45]. In this section, we present a principled way to construct quantizers that unifies previous efforts, removes tedious calculations, and also brings theoretical convenience when it comes down to analyzing the convergence of the algorithms in Section 2.

Our construction is based on the proximal map $\mathbf{P}_r^\mu: \mathbb{R}^d \rightrightarrows \mathbb{R}^d$ of a (closed) function r :

$$\mathbf{P}_r^\mu(\mathbf{w}^*) = \arg \min_{\mathbf{w}} \frac{1}{2\mu} \|\mathbf{w} - \mathbf{w}^*\|_2^2 + r(\mathbf{w}), \quad (8)$$

where $\mu > 0$ is a smoothing parameter. The proximal map is well-defined as long as the function r is lower bounded by a quadratic function, in particular, when r is bounded from below. If r is proper and (closed) convex, then the minimizer on the right-hand side of (8) is uniquely attained, while for general (nonconvex) functions the proximal map may be multi-valued (hence the notation \rightrightarrows). If $r = \iota_Q$ is the indicator function (see (15) below), then \mathbf{P}_r^μ reduces to the familiar projector \mathbf{P}_Q (for any μ). Remarkably, a complete characterization of such maps on the real line is available:

Theorem 3.1 ([Proposition 3, 48]). A (possibly multi-valued) map $P: \mathbb{R} \rightrightarrows \mathbb{R}$ is a proximal map (of some function r) iff it is (nonempty) compact-valued, monotone and has a closed graph. The underlying function r is unique (up to addition of constants) iff P is convex-valued, while r is convex iff P is nonexpansive (i.e. 1-Lipschitz continuous).

The sufficient and necessary conditions of Theorem 3.1 allow one to design proximal maps P directly, without needing to know the underlying function r at all (even though it is possible to integrate a version of r from P). The point is that, as far as quantization algorithms are concerned, having P is enough, and hence one is excused from the tedious calculations in deriving P from r as is typical in existing works.

For example, the (univariate) mirror maps constructed in Ajanthan et al. [Theorem 1, 1], such as \tanh , are proximal maps according to Theorem 3.1: In our notation, the (Mirror Descent) updates of Ajanthan et al. [1] are the same as (3) with $P = (\nabla\Phi)^{-1}$ for some mirror map Φ . Since Φ is taken to be strictly convex in Ajanthan et al. [1], it is easy to verify that P satisfies all conditions of Theorem 3.1.²

The next result allows us to employ *stochastic* quantizers, where in each iteration we randomly choose one of P_i to quantize the weights³. We may also apply different quantizers to different layers of a neural network.

Theorem 3.2. Let $P_i: \mathbb{R}^d \rightrightarrows \mathbb{R}^d, i = [k]$, be proximal maps. Then, the averaged map

$$P := \sum_{i=1}^k \alpha_i P_i, \quad \text{where } \alpha_i \geq 0, \quad \sum_{i=1}^k \alpha_i = 1, \quad (9)$$

is also a proximal map. Similarly, the product map

$$P := P_1 \times P_2 \times \cdots \times P_k, \quad \mathbf{w}^* = (\mathbf{w}_1^*, \dots, \mathbf{w}_k^*) \mapsto (P_1(\mathbf{w}_1^*), \dots, P_k(\mathbf{w}_k^*)) \quad (10)$$

is a proximal map (from \mathbb{R}^{dk} to \mathbb{R}^{dk}).

(The proof of Theorem 3.2 and all other omitted proofs can be found in Appendix A.)

Example 3.3. Let Q be a quantization set (e.g. $Q = \{-1, 0, 1\}$). Clearly, the identity map id and the projector P_Q are proximal maps. Therefore, their convex combination $P^\mu = \frac{\text{id} + \mu P_Q}{1 + \mu}, \mu \geq 0$, is also a proximal map, which is exactly the quantizer used in Yin et al. [45].

Lastly, we mention that it is intuitively desirable to have $P^\mu \rightarrow P_Q$ when μ increases indefinitely, so that the quantizer eventually settles on *bona fide* discrete values in Q . This is easily achieved by letting minimizers of the underlying function r , or the (possibly larger set of) fixed points of P^μ , approach Q . We give one such construction by exploiting Theorem 3.1 and Theorem 3.2.

3.1 General Piecewise Linear Quantizers

Let $Q_j = \{q_{j,k}\}_{k=1}^{b_j}$ (with $q_{j,1} \leq \cdots \leq q_{j,b_j}$) be the quantization set for the j -th weight group⁴ $\mathbf{w}_j^* \in \mathbb{R}^{d_j}$. Let $p_{j,k+1} := \frac{q_{j,k} + q_{j,k+1}}{2}, k \in [b_j - 1]$, be the middle points. We introduce two parameters $\rho, \varrho \geq 0$ and define

$$\text{horizontal shifts:} \quad q_{j,k}^- := p_{j,k} \vee (q_{j,k} - \rho), \quad q_{j,k}^+ := p_{j,k+1} \wedge (q_{j,k} + \rho), \quad (11)$$

$$\text{vertical shifts:} \quad p_{j,k+1}^- := q_{j,k} \vee (p_{j,k+1} - \varrho), \quad p_{j,k+1}^+ := q_{j,k+1} \wedge (p_{j,k+1} + \varrho), \quad (12)$$

with $q_{j,1}^- = q_{j,1}$ and $q_{j,b_j}^+ = q_{j,b_j}$. Then, we define L_ρ^ϱ as the piece-wise linear map (that simply connects the points by straight lines):

$$L_\rho^\varrho(w^*) := \begin{cases} q_{j,k}, & \text{if } q_{j,k}^- \leq w^* \leq q_{j,k}^+ \\ q_{j,k} + (w^* - q_{j,k}^+) \frac{p_{j,k+1}^- - q_{j,k}^-}{p_{j,k+1}^- - q_{j,k}^+}, & \text{if } q_{j,k}^+ \leq w^* < p_{j,k+1} \\ p_{j,k+1}^+ + (w^* - p_{j,k+1}) \frac{q_{j,k+1}^+ - p_{j,k+1}^-}{q_{j,k+1}^+ - p_{j,k+1}^-}, & \text{if } p_{j,k+1} < w^* \leq q_{j,k+1}^- \end{cases} \quad (13)$$

²The above reasoning hinges on Φ being univariate. More generally, if a multivariate mirror map Φ is 1-strongly convex (as is typical in Mirror Descent), then it follows from Moreau [Corollaire 10.c, 30] that $P = (\nabla\Phi)^{-1}$, being a nonexpansion, is again a proximal map (of some convex function).

³This form of stochasticity still leads to *deterministic networks*, and is therefore conceptually different from *probabilistic* (quantized) networks [35, 39].

⁴In this section, time subscripts are not needed. Instead, we use subscripts to indicate weight groups.




Figure 1: Different instantiations of the proximal map L_ρ^ϱ in (13) for $Q = \{-1, 0, 1\}$.

for all $w^* \in \mathbf{w}_j^*$. At the middle points, $L_\rho^\varrho(p_{j,k+1})$ may take any value within the two limits. Following the commonly used weight clipping in BC [11], we may set $L_\rho^\varrho(w^*) = q_{j,1}$ for $w^* < q_{j,1}$ and $L_\rho^\varrho(w^*) = q_{j,b_j}$ for $w^* > q_{j,b_j}$, however, other choices may also work well. The proximal map L_ρ^ϱ of an example ternary component is visualized in Figure 1 for different choices of ρ and ϱ .

The (horizontal) parameter ρ controls the discretization vicinity within which a continuous weight will be pulled *exactly* into the discrete set Q_j , while the (vertical) parameter ϱ controls the slope (i.e. expansiveness) of each piece. It follows at once from Theorem 3.1 that L_ρ^ϱ is indeed a proximal map. In particular, setting $\rho = 0$ (hence continuous weights are only discretized in the limit) and $\varrho = \frac{\mu}{2(1+\mu)}$ we recover Example 3.3 (assuming w.l.o.g. that $q_{j,k+1} - q_{j,k} \equiv 1$). On the other hand, setting $\rho = \varrho$ leads to a generalization of the (binary) quantizer in Bai et al. [5], which keeps the slope to the constant 1 (while introducing jumps at middle points) and happens to be the proximal map of the distance function to Q_j [5]. Of course, setting $\rho = \varrho = 0$ yields the identity map and allows us to skip quantizing certain weights (as is common practice), while letting $\rho, \varrho \rightarrow \infty$ recovers the projector P_{Q_j} .

Needless to say, we may adapt the quantization set Q_j and the parameters ϱ and ρ for different weight groups, creating a multitude of quantization schemes. By Theorem 3.2, the overall operator remains a proximal map.

4 Demystifying BinaryConnect (BC)

Bai et al. [5] observed that the updates of BC are formally the same as the dual averaging (DA) algorithm [31, 42], even though the latter algorithm was originally proposed and analyzed only for convex problems. A lesser known fact is that (regularized) dual averaging itself is a special case of the generalized conditional gradient (GCG) algorithm. In this section, we first present the aforementioned fact, refine the observation of Bai et al. [5], and set up the stage for generalizing BC.

4.1 Generalized Conditional Gradient is Primal-Dual

We first present the generalized conditional gradient (GCG) algorithm [8, 49] and point out its ability to solve simultaneously the primal and dual problems.

Let us consider the following “regularized” problem:

$$\min_{\mathbf{w} \in \mathbb{R}^d} f(\mathbf{w}) := \ell(\mathbf{w}) + r(\mathbf{w}), \quad (14)$$

where r is a general (nonconvex) regularizer. Setting r to the indicator function of Q , i.e.,

$$r(\mathbf{w}) = \iota_Q(\mathbf{w}) = \begin{cases} 0, & \text{if } \mathbf{w} \in Q \\ \infty, & \text{otherwise} \end{cases}, \quad (15)$$

reduces (14) to the original problem (1). As we will see, incorporating an arbitrary r does not add any complication but will allow us to immediately generalize BC.

Introducing the Fenchel conjugate function ℓ^* (resp. r^*) of ℓ (resp. r):

$$\ell^*(\mathbf{w}^*) := \sup_{\mathbf{w}} \langle \mathbf{w}, \mathbf{w}^* \rangle - \ell(\mathbf{w}), \quad (16)$$

which is always (closed) convex even when ℓ itself is nonconvex, we state the Fenchel–Rockafellar dual problem [38]:

$$\min_{\mathbf{w}^* \in \mathbb{R}^d} \ell^*(-\mathbf{w}^*) + r^*(\mathbf{w}^*), \quad (17)$$

which, unlike the original problem (14), is always a convex problem.

We apply the generalized conditional gradient algorithm [8, 49] to solving the dual problem⁵ (17): Given \mathbf{w}_t^* , we linearize the function r^* and solve

$$\mathbf{z}_t^* = \left[\arg \min_{\mathbf{w}^*} \ell^*(-\mathbf{w}^*) + \langle \mathbf{w}^*, \mathbf{w}_t \rangle \right] = -\nabla \ell^{**}(\mathbf{w}_t), \quad \mathbf{w}_t := \nabla r^*(\mathbf{w}_t^*), \quad (18)$$

where we used the fact that $(\nabla \ell^*)^{-1} = \nabla \ell^{**}$. Then, we take the convex combination

$$\mathbf{w}_{t+1}^* = (1 - \lambda_t) \mathbf{w}_t^* + \lambda_t \mathbf{z}_t^*, \quad \text{where } \lambda_t \in [0, 1]. \quad (19)$$

The following theorem extends Bach [Proposition 4.2, 4] and Yu [Theorem 4.6, 47] to any λ_t :

Theorem 4.1. Suppose r^* is L -smooth (i.e. ∇r^* is L -Lipschitz continuous), then for any \mathbf{w} :

$$\sum_{\tau=0}^t \frac{\lambda_\tau}{\pi_\tau} [(\ell^{**} + r^{**})(\mathbf{w}_\tau) - (\ell^{**} + r^{**})(\mathbf{w})] \leq (1 - \lambda_0) \Delta(\mathbf{w}, \mathbf{w}_0) + \sum_{\tau=0}^t \frac{\lambda_\tau^2}{2\pi_\tau} L \|\mathbf{w}_\tau^* - \mathbf{z}_\tau^*\|_2^2, \quad (20)$$

where $\mathbf{w}_t := \nabla r^*(\mathbf{w}_t^*)$, $\pi_t := \prod_{\tau=1}^t (1 - \lambda_\tau)$, and $\pi_0 := 1$. $\Delta(\mathbf{w}, \mathbf{w}_t) := r^{**}(\mathbf{w}) - r^{**}(\mathbf{w}_t) - \langle \mathbf{w} - \mathbf{w}_t, \mathbf{w}_t^* \rangle$ is the Bregman divergence induced by the convex function r^{**} .

While the convergence of \mathbf{w}_t^* to the minimum of (17) is well-known (see e.g. Yu et al. [49]), the above result also implies that a properly averaged iterate $\bar{\mathbf{w}}_t$ also converges to the minimum of the dual problem of (17):

Corollary 4.2. Let $\bar{\mathbf{w}}_t := \sum_{\tau=0}^t \Lambda_{t,\tau} \mathbf{w}_\tau$, where $\Lambda_{t,\tau} := \frac{\lambda_\tau}{\pi_\tau} / H_t$ and $H_t := \sum_{\tau=0}^t \frac{\lambda_\tau}{\pi_\tau}$. Then, we have for any \mathbf{w} :

$$(\ell^{**} + r^{**})(\bar{\mathbf{w}}_t) - (\ell^{**} + r^{**})(\mathbf{w}) \leq \frac{(1 - \lambda_0) \Delta(\mathbf{w}, \mathbf{w}_0)}{H_t} + \frac{L}{2} \sum_{\tau=0}^t \lambda_\tau \Lambda_{t,\tau} \|\mathbf{w}_\tau^* - \mathbf{z}_\tau^*\|_2^2. \quad (21)$$

Assuming $\{\mathbf{z}_\tau^*\}$ is bounded (e.g. when ℓ^{**} is Lipschitz continuous), the right-hand side of (21) diminishes if $\lambda_t \rightarrow 0$ and $\sum_t \lambda_t = \infty$. Setting $\lambda_t = \frac{1}{t+1}$ recovers ergodic averaging $\bar{\mathbf{w}}_t = \frac{1}{t+1} \sum_{\tau=0}^t \mathbf{w}_\tau$ for which the right-hand side of (21) diminishes at the rate⁶ $O(\log t/t)$; see Appendix A.2 for details.

Thus, GCG solves problem (17) and its dual simultaneously.

4.2 BC \subseteq DA \subseteq GCG

We are now in a position to reveal the relationships among BinaryConnect (BC), (regularized) dual averaging (DA) and the generalized conditional gradient (GCG). Since Theorem 4.1 requires r^* to be L -smooth, in the event that it is not we resort to a smooth approximation known as the Moreau envelope [30]:

$$\mathcal{M}_{r^*}^\mu(\mathbf{w}^*) = \min_{\mathbf{z}^*} \frac{1}{2\mu} \|\mathbf{w}^* - \mathbf{z}^*\|_2^2 + r^*(\mathbf{z}^*), \quad (22)$$

where the minimizer is (by definition) exactly $\mathbf{P}_{r^*}^\mu(\mathbf{w}^*)$. It is well-known that $\mathcal{M}_{r^*}^\mu$ is $(1/\mu)$ -smooth and $(\mathcal{M}_{r^*}^\mu)^* = r^{**} + \frac{\mu}{2} \|\cdot\|_2^2$ [30]. We then apply GCG to the approximate dual problem:

$$\min_{\mathbf{w}^* \in \mathbb{R}^d} \ell^*(-\mathbf{w}^*) + \mathcal{M}_{r^*}^\mu(\mathbf{w}^*), \quad (23)$$

whose own Fenchel–Rockfellar dual is:

$$\left[\min_{\mathbf{w} \in \mathbb{R}^d} \ell^{**}(\mathbf{w}) + (\mathcal{M}_{r^*}^\mu)^*(\mathbf{w}) \right] = \min_{\mathbf{w} \in \mathbb{R}^d} \ell^{**}(\mathbf{w}) + r^{**}(\mathbf{w}) + \frac{\mu}{2} \|\mathbf{w}\|_2^2. \quad (24)$$

⁵GCG is usually applied to solving the primal problem (14) directly. Our choice of the dual problem here is to facilitate later comparison with dual averaging and BinaryConnect.

⁶The log factor can be removed by setting, for example, $\lambda_t = \frac{2}{2+t}$ instead.

The updates of GCG applied to the approximate problem (23) are thus:

$$\begin{aligned}\mathbf{w}_t &:= \nabla \mathcal{M}_{r^{**}}^\mu(\mathbf{w}_t^*) = \mathbf{P}_{r^{**}}^{1/\mu}(\mathbf{w}_t^*/\mu) \quad (\text{see Proposition A.2 for derivation}) & (25) \\ \mathbf{w}_{t+1}^* &= (1 - \lambda_t)\mathbf{w}_t^* + \lambda_t \mathbf{z}_t^*, \quad \text{where } \mathbf{z}_t^* = -\nabla \ell^{**}(\mathbf{w}_t), & (26)\end{aligned}$$

which is exactly the updates of (regularized) dual averaging [42] for convex problems where⁷ $r^{**} = r$ and $\ell^{**} = \ell$. Nesterov [31] motivated dual averaging by the natural desire of non-decreasing step sizes, whereas conventional subgradient algorithms “counter-intuitively” assign smaller step sizes to more recent iterates instead. Based on our explanation, we conclude this is possible because dual averaging solves an (approximate) smoothed dual problem, hence we can afford to use a constant (rather than a diminishing/decreasing) step size.

Defining $\pi_t := \prod_{s=1}^t (1 - \lambda_s)$ for $t \geq 1$, $\pi_0 := 1$, $\pi_{-1} := (1 - \lambda_0)^{-1}$, and setting $\mathbf{w}_t^* = \mathbf{w}_t^*/\pi_{t-1}$, we have:

$$\mathbf{w}_t = \mathbf{P}_{r^{**}}^{1/\mu}(\pi_{t-1}\mathbf{w}_t^*/\mu), \quad \mathbf{w}_{t+1}^* = \mathbf{w}_t^* - \frac{\lambda_t}{\pi_t} \nabla \ell^{**}(\mathbf{w}_t). \quad (27)$$

Let us now reparameterize

$$\eta_t := \frac{\lambda_t}{\pi_t} \implies \lambda_t = \frac{\eta_t}{1 + \sum_{\tau=1}^t \eta_\tau} \text{ and } \frac{1}{\pi_t} = 1 + \sum_{\tau=1}^t \eta_\tau, \quad (28)$$

for $t \geq 1$. If we also allow $\mu = \mu_t$ to change adaptively from iteration to iteration (as in dual averaging), in particular, if $\mu_t = \pi_{t-1}$, we obtain the familiar update:

$$\mathbf{w}_t = \mathbf{P}_{r^{**}}^{1/\pi_{t-1}}(\mathbf{w}_t^*), \quad \mathbf{w}_{t+1}^* = \mathbf{w}_t^* - \eta_t \nabla \ell^{**}(\mathbf{w}_t). \quad (29)$$

For nonconvex problems, we may replace r^{**} and ℓ^{**} with their nonconvex counterparts r and ℓ , respectively. We remark that r^{**} (resp. ℓ^{**}) is the largest convex function that is (pointwise) majorized by r (resp. ℓ). In particular, with $r = \iota_Q$ and $\mathbf{P}_r^{1/\mu} = \mathbf{P}_Q$ (for any μ) we recover the BinaryConnect update (3). While the parameter μ plays no role when r is an indicator function, we emphasize that for general r we should use the quantizer $\mathbf{P}_r^{1/\pi_{t-1}}$, where importantly $1/\pi_{t-1} \rightarrow \infty$ hence the quantizer converges to minimizers of r asymptotically. Neglecting this crucial detail may lead to suboptimality as is demonstrated in the following example:

Example 4.3. Bai et al. [5] constructed the following intriguing example:

$$\ell(w) = \frac{1}{2}w^2, \quad Q = \{\pm 1\}, \quad \mathbf{P}_r^{1/\mu}(w) = \text{sign}(w) \frac{\epsilon|w| + \frac{1}{\mu}}{\epsilon + \frac{1}{\mu}} \text{ for } |w| \leq 1, \quad (30)$$

and they showed non-convergence of the algorithm $w \leftarrow w - \eta \nabla \ell(\mathbf{P}_r^{1/\mu}(w))$, where μ is a fixed constant. If we use $\mathbf{P}_r^{1/\pi_{t-1}}$ with some diverging $1/\pi_{t-1}$ instead, the resulting BinaryConnect, with diminishing η_t or ergodic averaging, would actually converge to 0 (since $\mathbf{P}_r^{1/\pi_{t-1}} \rightarrow \text{sign}$).

5 ProxConnect (PC): A Generalization of BinaryConnect

We are now ready to generalize BC by combining the results from Section 3 and Section 4. Replacing the convex envelopes, ℓ^{**} and r^{**} , with their nonconvex counterparts and replacing deterministic gradients with stochastic gradients (as well as the change-of-variable $\mathbf{w}_t^* \rightarrow \tilde{\mathbf{w}}_t^*$), we obtain from (29) a *family* of algorithms which we term ProxConnect (PC):

$$\mathbf{w}_t = \mathbf{P}_r^{1/\pi_{t-1}}(\tilde{\mathbf{w}}_t^*), \quad \tilde{\mathbf{w}}_{t+1}^* = \tilde{\mathbf{w}}_t^* - \eta_t \tilde{\nabla} \ell(\tilde{\mathbf{w}}_t), \quad (31)$$

where the quantizer $\mathbf{P}_r^{1/\pi_{t-1}}$ may be designed directly by following Section 3. We have already seen in Section 4 that BC belongs to PC by choosing $\mathbf{P}_r^{1/\pi_{t-1}} = \mathbf{P}_Q$ (in which case π_{t-1} plays no role).

The analysis in Section 4, initially tailored to convex functions, immediately generalizes to the nonconvex algorithm PC (for nonconvex ℓ , nonconvex r , and stochastic gradients $\tilde{\nabla} \ell$):

⁷That is, if we set $\lambda_t = 1/t$ and allow for time-dependent $\mu_t = t$; see Xiao [Algorithm 1, 42].

Theorem 5.1. Fix any \mathbf{w} , the iterates in (31) satisfy:

$$\sum_{\tau=s}^t \eta_\tau [\langle \mathbf{w}_\tau - \mathbf{w}, \tilde{\nabla} \ell(\mathbf{w}_\tau) \rangle + r(\mathbf{w}_\tau) - r(\mathbf{w})] \leq \Delta_{s-1}(\mathbf{w}) - \Delta_t(\mathbf{w}) + \sum_{\tau=s}^t \Delta_\tau(\mathbf{w}_\tau), \quad (32)$$

where $\Delta_\tau(\mathbf{w}) := r_\tau(\mathbf{w}) - r_\tau(\mathbf{w}_{\tau+1}) - \langle \mathbf{w} - \mathbf{w}_{\tau+1}, \mathbf{w}_{\tau+1}^* \rangle$ is the Bregman divergence induced by the (possibly nonconvex) function $r_\tau(\mathbf{w}) := \frac{1}{\pi_\tau} r(\mathbf{w}) + \frac{1}{2} \|\mathbf{w}\|_2^2$.

The summand on the left-hand side of (32) is related to the duality gap in Yu et al. [49], which is a natural measure of stationarity for the nonconvex problem (14). Indeed, it reduces to the familiar ones when convexity is present:

Corollary 5.2. For convex ℓ and any \mathbf{w} , the iterates in (31) satisfy:

$$\min_{\tau=s, \dots, t} \mathbf{E}[f(\mathbf{w}_\tau) - f(\mathbf{w})] \leq \frac{1}{\sum_{\tau=s}^t \eta_\tau} \cdot \mathbf{E}[\Delta_{s-1}(\mathbf{w}) - \Delta_t(\mathbf{w}) + \sum_{\tau=s}^t \Delta_\tau(\mathbf{w}_\tau)]. \quad (33)$$

If r is also convex, then

$$\min_{\tau=s, \dots, t} \mathbf{E}[f(\mathbf{w}_\tau) - f(\mathbf{w})] \leq \frac{1}{\sum_{\tau=s}^t \eta_\tau} \cdot \mathbf{E}[\Delta_{s-1}(\mathbf{w}) + \sum_{\tau=s}^t \frac{\eta_\tau^2}{2} \|\tilde{\nabla} \ell(\mathbf{w}_\tau)\|_2^2], \quad (34)$$

and

$$\mathbf{E}[f(\bar{\mathbf{w}}_t) - f(\mathbf{w})] \leq \frac{1}{\sum_{\tau=s}^t \eta_\tau} \cdot \mathbf{E}[\Delta_{s-1}(\mathbf{w}) + \sum_{\tau=s}^t \frac{\eta_\tau^2}{2} \|\tilde{\nabla} \ell(\mathbf{w}_\tau)\|_2^2], \quad (35)$$

where $\bar{\mathbf{w}}_t = \frac{\sum_{\tau=s}^t \eta_\tau \mathbf{w}_\tau}{\sum_{\tau=s}^t \eta_\tau}$.

The right-hand sides of (34) and (35) diminish iff $\eta_t \rightarrow 0$ and $\sum_t \eta_t = \infty$ (assuming boundedness of the stochastic gradient). We note some trade-off in choosing the step size η_τ : both the numerator and denominator of the right-hand sides of (34) and (35) are increasing w.r.t. η_τ . The same conclusion can be drawn for (33) and (32), where Δ_τ also depends on η_τ (through the accumulated magnitude of $\mathbf{w}_{\tau+1}^*$). A detailed analysis may need to take specific properties of r or P into account [45].

ProxQuant vs ProxConnect. It is worthwhile to point out one important difference between ProxQuant and ProxConnect: Bai et al. [5] proved convergence (to some notion of stationarity) of ProxQuant for a fixed quantizer (see Bai et al. [Theorem 5.1, 5]), i.e., P_r^μ for a fixed μ , but their experiments relied on incrementing μ so that their quantizer approaches the projector P_Q . This creates some discrepancy between theory and practice. The same comment also applies to [1]. In contrast, ProxConnect is derived from a rigorous theory that automatically justifies a diverging μ . In particular, choosing a constant step size $\eta_\tau \equiv \eta_0$ would lead to $1/\pi_{t-1} \propto t$, matching the current practice that is now justifiable if r is strongly convex; see Appendix A.3.1 for details.

Connection to Existing Algorithms. Besides the obvious BC, ProxConnect generalizes many other quantization algorithms. As it turns out, many of these algorithms can also be realized using our proposed proximal quantizer L_ρ^θ from Section 3; see Table 1 for a sample summary.

reverseProxConnect. In Section 2, we discussed the idea of reversing BinaryConnect. As ProxConnect generalizes BC, we also present reverseProxConnect (rPC) as a generalization of reversing BinaryConnect:

$$\mathbf{w}_t = P_r^{1/\pi_{t-1}}(\mathbf{w}_t^*), \quad \mathbf{w}_{t+1}^* = \mathbf{w}_t - \eta_t \tilde{\nabla} \ell(\mathbf{w}_t^*), \quad (36)$$

In contrast to reversing BinaryConnect, rPC is not completely without merit: it evaluates the gradient at the continuous weights \mathbf{w}_t^* and hence is able to exploit a richer landscape of the loss. Even when stuck at a fixed discrete weight \mathbf{w}_t , rPC may still accumulate sizable updates (as long as the step size and the gradient remain sufficiently large) to allow it to eventually jump out of \mathbf{w}_t : note that the continuous weights \mathbf{w}_t^* still get updated. Finally, for constant step size η and π , we note that fixed points of rPC, when existent, satisfy:

$$\mathbf{w}^* = P_r^{1/\pi}(\mathbf{w}^*) - \eta \nabla \ell(\mathbf{w}^*) \iff \mathbf{w}^* = (\text{id} + \eta \nabla \ell)^{-1} \left(\text{id} + \eta \left[\frac{\partial r}{\partial \eta} \right] \right)^{-1} (\mathbf{w}^*) \quad (37)$$

$$=: \mathcal{B} \left(\eta, \nabla \ell, \frac{\partial r}{\partial \eta} \right) (\mathbf{w}^*), \quad (38)$$

Table 1: A sample summary of existing quantization algorithms. The PC column indicates if the method is a special case of our proposed ProxConnect algorithm. The L_ρ^ϱ -column indicates if the method uses a quantizer which is a special case of our general quantizer L_ρ^ϱ introduced in Section 3. If so, the ρ, ϱ -column states how ρ and ϱ were chosen (in practice): increasing \nearrow , fixed to 0, or fixed to ∞ . Other than ProxQuant-Ternary, all methods can compute their quantizers in a single neural network pass. †: TrainedTernary methods might use a quantizer different than L_ρ^ϱ to improve performance.

Method	PC	One-pass	Learnable parameters	L_ρ^ϱ	ρ, ϱ
ProxQuant-Binary- L_1 [5]	✗	✓	✗	✓	\nearrow, \nearrow
ProxQuant-Ternary [5]	✗	✗	✗	✗	-
BinaryConnect [11]	✓	✓	✗	✓	∞, ∞
BinaryRelax [45]	✓	✓	✗	✓	0, \nearrow
BinaryWeight [37]	✓	✓	✓	✗	-
MirrorDescentView [1]	✓	✓	✗	✗	-
TrainedTernary† [53]	✓	✓	✓	✓	∞, ∞
TernaryWeight [26]	✓	✓	✓	✗	-
ProxConnect (ours)	-	✓	✓	✓	\nearrow, \nearrow

where for simplicity we assumed deterministic gradients $\nabla \ell$ and r to be convex such that $P_r^{1/\pi} = (\text{id} + \partial r/\pi)^{-1}$. The operator \mathcal{B} is known as the backward-backward update (as opposed to the forward-backward update in ProxQuant), and it is known that when $\eta \rightarrow 0$ slowly, backward-backward updates converge to a stationary point [34]. Thus, despite of our current limited understanding of rPC, there is some reason (and empirical evidence as shown in Section 6) to believe it might still be interesting.

Importance of GCG framework: Deriving PC from the GCG framework may let us transfer recent advances on GCG [7, 14, 50] to neural network quantization. Furthermore, it is the cornerstone in justifying the widely-adopted diverging smoothing parameter.

6 Experiments

6.1 Classification on CIFAR-10

We perform image classification on CIFAR-10 [24] using ResNet20 and ResNet56 [19], comparing BinaryConnect [11] with ProxQuant [5] and (reverse)ProxConnect using our proposed proximal operator L_ρ^ϱ . For fair comparison, we set $\rho = \varrho$ in L_ρ^ϱ as this resembles the quantizer (for binary quantization) used in the original ProxQuant algorithm. Similar to Bai et al. [5], we increase the parameter ρ (or equivalently ϱ) linearly: $\rho_t = (1 + t/B)\rho_0$. In contrast to Bai et al. [5], however, we increase ρ after every gradient step rather than after every epoch as this is more in line with our analysis. We treat ρ_0 as a hyperparameter for which we conduct a small grid search. We consider binary ($Q = \{-1, 1\}$), ternary ($Q = \{-1, 0, 1\}$) and quaternary ($Q = \{-1, -0.3, 0.3, 1\}$) quantization. Details for all CIFAR-10 experiments can be found in Appendix B.1.

Fine-Tuning Pretrained Models. In this experiment, all quantization algorithms are initialized with a pretrained ResNet. The test accuracies of the pretrained ResNet20 and ResNet56 are 92.01 and 93.01, respectively. Table 2 shows the final test accuracies for the different models. For ProxQuant and (reverse)ProxConnect we respectively picked the best ρ_0 values; results for all ρ_0 values can be found in Appendix B.1.7. (Reverse)ProxConnect outperforms the other two methods on all six settings.

End-To-End Training. To save computational costs, it is important that quantization algorithms also perform well when they are not initialized with pretrained full-precision model. We therefore compare the four methods for randomly initialized models; see Table 3 for the results. ProxConnect outperforms all other methods on all six tasks. Interestingly, ProxQuant and reverseProxConnect perform considerably worse for all six tasks when compared to fine-tuning. The performance drop of BinaryConnect and ProxConnect when compared to fine-tuning is only significant for ternary

Table 2: Fine-tuning pretrained ResNets. Final test accuracy: mean and standard deviation in 3 runs.

Model	Quantization	BC [11]	PQ [5]	rPC (ours)	PC (ours)
ResNet20	Binary	90.31 (0.00)	89.94 (0.10)	89.98 (0.17)	90.31 (0.21)
	Ternary	74.95 (0.16)	91.46 (0.06)	91.47 (0.19)	91.37 (0.18)
	Quaternary	91.43 (0.07)	91.43 (0.21)	91.43 (0.06)	91.81 (0.14)
ResNet56	Binary	92.22 (0.12)	92.33 (0.06)	92.47 (0.29)	92.65 (0.16)
	Ternary	74.68 (1.4)	93.07 (0.02)	92.84 (0.11)	93.25 (0.12)
	Quaternary	93.20 (0.06)	92.82 (0.16)	92.91 (0.26)	93.42 (0.12)

Table 3: End-to-end training of ResNets. Final test accuracy: mean and standard deviation in 3 runs.

Model	Quantization	BC [11]	PQ [5]	rPC (ours)	PC (ours)
ResNet20	Binary	87.51 (0.21)	81.59 (0.75)	81.82 (0.32)	89.92 (0.65)
	Ternary	27.10 (0.21)	47.98 (1.30)	47.17 (1.94)	84.09 (0.16)
	Quaternary	89.91 (0.09)	85.29 (0.09)	85.05 (0.27)	90.17 (0.14)
ResNet56	Binary	89.79 (0.45)	86.13 (1.71)	86.25 (1.50)	91.26 (0.59)
	Ternary	30.31 (7.79)	50.54 (3.68)	42.95 (1.57)	84.36 (0.75)
	Quaternary	90.69 (0.57)	87.81 (1.60)	87.30 (1.02)	91.70 (0.14)

quantization. We found that ProxQuant and reverseProxConnect can be quite sensitive to the choice of ρ_0 , whereas ProxConnect is stable in this regard; see Appendix B.1.7.

6.2 Classification on ImageNet

We perform a small study on ImageNet [12] using ResNet18 [19]. As can be seen in Table 4, BC performs slightly better for fine-tuning whereas PC performs slightly better for end-to-end training. This is not an exhaustive study, but rather a first indication that PC can yield competitive performance on large scale datasets. For more details on the experiments see Appendix B.2.

Table 4: Fine-tuning (left) and end-to-end training (right). Final test accuracy: mean and standard deviation over three runs.

BC [11]	PC (ours)		BC [11]	PC (ours)	
	$\rho_0 = 2e-2$	$\rho_0 = 4e-2$		$\rho_0 = 2.5e-3$	$\rho_0 = 5e-3$
65.84 (0.04)	65.44 (0.13)	65.70 (0.04)	63.79 (0.12)	63.89 (0.14)	63.67 (0.12)

7 Conclusion

Capitalizing on a principled approach for designing quantizers and a surprising connection between BinaryConnect and the generalized conditional gradient (GCG) algorithm, we proposed ProxConnect as a unification and generalization of existing neural network quantization algorithms. Our analysis refines prior convergence guarantees and our experiments confirm the competitiveness of ProxConnect. In future work, we plan to apply ProxConnect to training other models such as transformers. The connection with GCG also opens the possibility for further acceleration.

Acknowledgments and Disclosure of Funding

We thank the anonymous reviewers for their constructive comments as well as the area chair and the senior area chair for overseeing the review process. YY thanks NSERC for funding support. We thank Jingjing Wang for some early discussions.

References

- [1] Thalaiyasingam Ajanthan, Kartik Gupta, Philip Torr, Richard Hartley, and Puneet Dokania. [Mirror Descent View for Neural Network Quantization](#). In *International Conference on Artificial Intelligence and Statistics*, 2021.
- [2] Milad Alizadeh, Javier Fernández-Marqués, Nicholas D. Lane, and Yarin Gal. [An Empirical study of Binary Neural Networks’ Optimisation](#). In *International Conference on Learning Representations*, 2019.
- [3] Alexander G. Anderson and Cory P. Berg. [The High-Dimensional Geometry of Binary Neural Networks](#). In *International Conference on Learning Representations*, 2018.
- [4] Francis Bach. [Duality Between Subgradient and Conditional Gradient Methods](#). *SIAM Journal on Optimization*, 25(1):115–129, 2015.
- [5] Yu Bai, Yu-Xiang Wang, and Edo Liberty. [ProxQuant: Quantized Neural Networks via Proximal Operators](#). In *International Conference on Learning Representations*, 2018.
- [6] Yoshua Bengio, Nicholas Léonard, and Aaron Courville. [Estimating or Propagating Gradients Through Stochastic Neurons for Conditional Computation](#). *arXiv:1308.3432*, 2013.
- [7] Leonard Berrada, Andrew Zisserman, and M. Pawan Kumar. [Deep Frank–Wolfe For Neural Network Optimization](#). In *International Conference on Learning Representations*, 2019.
- [8] Kristian Bredies and Dirk A. Lorenz. [Iterated Hard Shrinkage for Minimization Problems with Sparsity Constraints](#). *SIAM Journal on Scientific Computing*, 30(2):657–683, 2008.
- [9] Tom Brown and et al. [Language Models are Few-Shot Learners](#). In *Advances in Neural Information Processing Systems*, pages 1877–1901, 2020.
- [10] Shangyu Chen, Wenya Wang, and Sinno Jialin Pan. [MetaQuant: Learning to Quantize by Learning to Penetrate Non-differentiable Quantization](#). In *Advances in Neural Information Processing Systems*, 2019.
- [11] Matthieu Courbariaux, Yoshua Bengio, and Jean-Pierre David. [BinaryConnect: Training Deep Neural Networks with binary weights during propagations](#). In *Advances in Neural Information Processing Systems*, volume 28, pages 3123–3131, 2015.
- [12] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. [ImageNet: A large-scale hierarchical image database](#). In *IEEE Conference on Computer Vision and Pattern Recognition*, pages 248–255, 2009.
- [13] Yukun Ding, Jinglan Liu, Jinjun Xiong, and Yiyu Shi. [On the Universal Approximability and Complexity Bounds of Quantized ReLU Neural Networks](#). In *International Conference on Learning Representations*, 2019.
- [14] Pavel Dvurechensky, Petr Ostroukhov, Kamil Safin, Shimrit Shtern, and Mathias Staudigl. [Self-Concordant Analysis of Frank–Wolfe Algorithms](#). In *International Conference on Machine Learning*, 2020.
- [15] Yunhui Guo. [A Survey on Methods and Theories of Quantized Neural Networks](#). *arXiv:1808.04752*, 2018.
- [16] Suyog Gupta, Ankur Agrawal, Kailash Gopalakrishnan, and Pritish Narayanan. [Deep Learning with Limited Numerical Precision](#). In *International Conference on Machine Learning*, 2015.
- [17] Kai Han, Yunhe Wang, Yixing Xu, Chunjing Xu, Enhua Wu, and Chang Xu. [Training Binary Neural Networks through Learning with Noisy Supervision](#). In *International Conference on Machine Learning*, 2020.
- [18] Song Han, Xingyu Liu, Huizi Mao, Jing Pu, Ardavan Pedram, Mark A Horowitz, and William J Dally. [EIE: Efficient Inference Engine on Compressed Deep Neural Network](#). *ACM SIGARCH Computer Architecture News*, 44(3):243–254, 2016.

- [19] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. [Deep Residual Learning for Image Recognition](#). In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2016.
- [20] Koen Helwegen, James Widdicombe, Lukas Geiger, Zechun Liu, Kwang-Ting Cheng, and Roeland Nusselder. [Latent Weights Do Not Exist: Rethinking Binarized Neural Network Optimization](#). In *Advances in Neural Information Processing Systems*, 2019.
- [21] Sara Hooker, Nyalleng Moorosi, Gregory Clark, Samy Bengio, and Emily Denton. [Characterising Bias in Compressed Models](#). *arXiv:2010.03058*, 2020.
- [22] Lu Hou, Ruiliang Zhang, and James T. Kwok. [Analysis of Quantized Models](#). In *International Conference on Learning Representations*, 2019.
- [23] Itay Hubara, Matthieu Courbariaux, Daniel Soudry, Ran El-Yaniv, and Yoshua Bengio. [Quantized Neural Networks: Training Neural Networks with Low Precision Weights and Activations](#). *Journal of Machine Learning Research*, 18(1):6869–6898, 2017.
- [24] Alex Krizhevsky. [Learning Multiple Layers of Features from Tiny Images](#), 2009. The MIT License.
- [25] Cong Leng, Zesheng Dou, Hao Li, Shenghuo Zhu, and Rong Jin. [Extremely Low Bit Neural Network: Squeeze the Last Bit Out With ADMM](#). In *Proceedings of the Thirty-Second AAAI Conference on Artificial Intelligence*, 2018.
- [26] Fengfu Li, Bo Zhang, and Bin Liu. [Ternary Weight Networks](#). *arXiv:1605.04711*, 2016.
- [27] Hao Li, Soham De, Zheng Xu, Christoph Studer, Hanan Samet, and Tom Goldstein. [Training Quantized Nets: A Deeper Understanding](#). In *Advances in Neural Information Processing Systems*, 2017.
- [28] Ji Lin, Wei-Ming Chen, Yujun Lin, Chuang Gan, and Song Han. [MCUNet: Tiny Deep Learning on IoT Devices](#). *Advances in Neural Information Processing Systems*, 2020.
- [29] Brais Martinez, Jing Yang, Adrian Bulat, and Georgios Tzimiropoulos. [Training binary neural networks with real-to-binary convolutions](#). In *International Conference on Learning Representations*, 2020.
- [30] Jean-Jacques Moreau. [Proximité et dualité dans un espace hilbertien](#). *Bulletin de la Société Mathématique de France*, 93:273–299, 1965.
- [31] Yurii Nesterov. [Primal-dual subgradient methods for convex problems](#). *Mathematical Programming*, 120(1):221–259, 2009.
- [32] E. Park, J. Ahn, and S. Yoo. [Weighted-Entropy-Based Quantization for Deep Neural Networks](#). In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 7197–7205, 2017.
- [33] Vahid Partovi Nia and Mouloud Belbahri. [Binary Quantizer](#). *Journal of Computational Vision and Imaging Systems*, 4(1):3, 2018.
- [34] Gregory B Passty. [Ergodic Convergence to a Zero of the Sum of Monotone Operators in Hilbert Space](#). *Journal of Mathematical Analysis and Applications*, 72(2):383–390, 1979.
- [35] Jorn WT Peters and Max Welling. [Probabilistic Binary Neural Networks](#). *arXiv:1809.03368*, 2018.
- [36] Haotong Qin, Ruihao Gong, Xianglong Liu, Xiao Bai, Jingkuan Song, and Nicu Sebe. [Binary Neural Networks: A Survey](#). *Pattern Recognition*, 105:107281, 2020.
- [37] Mohammad Rastegari, Vicente Ordonez, Joseph Redmon, and Ali Farhadi. [XNOR-Net: ImageNet Classification Using Binary Convolutional Neural Networks](#). In *European Conference on Computer Vision*, 2016.
- [38] R. Tyrrell Rockafellar and Roger J-B Wets. [Variational Analysis](#). Springer, 1998.

- [39] Oran Shayer, Dan Levi, and Ethan Fetaya. [Learning Discrete Weights Using the Local Reparameterization Trick](#). In *International Conference on Learning Representations*, 2018.
- [40] Emma Strubell, Ananya Ganesh, and Andrew McCallum. [Energy and Policy Considerations for Deep Learning in NLP](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 3645–3650, 2019.
- [41] Yanzhi Wang, Zheng Zhan, Liang Zhao, Jian Tang, Siyue Wang, Jiayu Li, Bo Yuan, Wujie Wen, and Xue Lin. [Universal Approximation Property and Equivalence of Stochastic Computing-Based Neural Networks and Binary Neural Networks](#). In *Proceedings of the Thirty-Third AAAI Conference on Artificial Intelligence*, 2019.
- [42] Lin Xiao. [Dual Averaging Methods for Regularized Stochastic Learning and Online Optimization](#). *Journal of Machine Learning Research*, 11:2543–2596, 2010.
- [43] Chen Xu, Jianqiang Yao, Zhouchen Lin, Wenwu Ou, Yuanbin Cao, Zhirong Wang, and Hongbin Zha. [Alternating Multi-bit Quantization for Recurrent Neural Networks](#). In *International Conference on Learning Representations*, 2018.
- [44] P Yin, S Zhang, YY Qi, and J Xin. [Quantization and Training of Low Bit-Width Convolutional Neural Networks for Object Detection](#). *Journal of Computational Mathematics*, 37(3), 2019.
- [45] 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](#). *SIAM Journal on Imaging Sciences*, 11(4):2205–2223, 2018.
- [46] Penghang Yin, Jiancheng Lyu, Shuai Zhang, Stanley J. Osher, Yingyong Qi, and Jack Xin. [Understanding Straight-Through Estimator in Training Activation Quantized Neural Nets](#). In *International Conference on Learning Representations*, 2019.
- [47] Yaoliang Yu. [Fast Gradient Methods for Structured Sparsity](#). PhD thesis, University of Alberta, 2013.
- [48] Yaoliang Yu, Xun Zheng, Micol Marchetti-Bowick, and Eric Xing. [Minimizing Nonconvex Non-Separable Functions](#). In *Artificial Intelligence and Statistics*, 2015.
- [49] Yaoliang Yu, Xinhua Zhang, and Dale Schuurmans. [Generalized Conditional Gradient for Structured Sparse Estimation](#). *Journal of Machine Learning Research*, 18:1–46, 2017.
- [50] Mingrui Zhang, Lin Chen, Aryan Mokhtari, Hamed Hassani, and Amin Karbasi. [Quantized Frank–Wolfe: Faster Optimization, Lower Communication, and Projection Free](#). In *International Conference on Artificial Intelligence and Statistics*, 2020.
- [51] Aojun Zhou, Anbang Yao, Yiwen Guo, Lin Xu, and Yurong Chen. [Incremental Network Quantization: Towards Lossless CNNs with Low-precision Weights](#). In *International Conference on Learning Representations*, 2017.
- [52] Shuchang Zhou, Yuxin Wu, Zekun Ni, Xinyu Zhou, He Wen, and Yuheng Zou. [DoReFa-Net: Training low Bitwidth Convolutional Neural Networks with Low Bitwidth Gradients](#). *arXiv:1606.06160*, 2016.
- [53] Chenzhuo Zhu, Song Han, Huizi Mao, and William J Dally. [Trained Ternary Quantization](#). In *International Conference on Learning Representations*, 2018.

Checklist

1. For all authors...
 - (a) Do the main claims made in the abstract and introduction accurately reflect the paper’s contributions and scope? [Yes]
 - (b) Did you describe the limitations of your work? [Yes] As discussed, the goal of our paper is to improve the understanding of BinaryConnect [11] and obtain novel generalizations. We therefore focus more on theory than empirical results. As a result, we only did extensive experiments on one dataset (CIFAR-10) with one architecture (ResNet20/56). As discussed in Section 6.2, our experiments on ImageNet are not exhaustive; both due to environmental reasons as well as limited availability of compute. In the future, we aim to do more exhaustive testing of ProxConnect (different architectures and problems), while keeping in mind the carbon footprint that comes with it.
 - (c) Did you discuss any potential negative societal impacts of your work? [Yes] See the first paragraph of Section 1: quantization leads to decreased carbon footprint for training and inference (positive societal impact), however, it may come at the cost of increased bias (negative societal impact).
 - (d) Have you read the ethics review guidelines and ensured that your paper conforms to them? [Yes]
2. If you are including theoretical results...
 - (a) Did you state the full set of assumptions of all theoretical results? [Yes]
 - (b) Did you include complete proofs of all theoretical results? [Yes] All proofs are provided in the Appendix.
3. If you ran experiments...
 - (a) Did you include the code, data, and instructions needed to reproduce the main experimental results (either in the supplemental material or as a URL)? [No] We will provide the code once it has been approved by our internal process. We only ran experiments on publicly available datasets.
 - (b) Did you specify all the training details (e.g., data splits, hyperparameters, how they were chosen)? [Yes] See Section 6, Appendix B.1 and Appendix B.2.
 - (c) Did you report error bars (e.g., with respect to the random seed after running experiments multiple times)? [Yes] All experiments are over three random seeds; we provide mean and standard deviation for all results.
 - (d) Did you include the total amount of compute and the type of resources used (e.g., type of GPUs, internal cluster, or cloud provider)? [Yes] The amount of compute in the exploration stage of our project was insignificant compared to the experiments reported in this work. See Appendix B.1.6 and Appendix B.8 for the total amount of compute we used for the experiments reported in this work.
4. If you are using existing assets (e.g., code, data, models) or curating/releasing new assets...
 - (a) If your work uses existing assets, did you cite the creators? [Yes] See Section 6.1 and Section 6.2.
 - (b) Did you mention the license of the assets? [Yes] Yes, the license of CIFAR-10 (MIT license) is included in the bibliography entry. The ImageNet dataset is “is available for free to researchers for non-commercial use”.
 - (c) Did you include any new assets either in the supplemental material or as a URL? [No] We will provide the code once it has been approved by our internal process.
 - (d) Did you discuss whether and how consent was obtained from people whose data you’re using/curating? [N/A]
 - (e) Did you discuss whether the data you are using/curating contains personally identifiable information or offensive content? [N/A]
5. If you used crowdsourcing or conducted research with human subjects...
 - (a) Did you include the full text of instructions given to participants and screenshots, if applicable? [N/A]
 - (b) Did you describe any potential participant risks, with links to Institutional Review Board (IRB) approvals, if applicable? [N/A]

(c) Did you include the estimated hourly wage paid to participants and the total amount spent on participant compensation? [N/A]

A Proofs

In this section we include all proofs omitted in the main paper, and supply some additional comments.

A.1 Proofs for Section 3

Theorem 3.2. Let $P_i : \mathbb{R}^d \rightrightarrows \mathbb{R}^d, i = [k]$, be proximal maps. Then, the averaged map

$$P := \sum_{i=1}^k \alpha_i P_i, \quad \text{where } \alpha_i \geq 0, \quad \sum_{i=1}^k \alpha_i = 1, \quad (9)$$

is also a proximal map. Similarly, the product map

$$P := P_1 \times P_2 \times \cdots \times P_k, \quad \mathbf{w}^* = (\mathbf{w}_1^*, \dots, \mathbf{w}_k^*) \mapsto (P_1(\mathbf{w}_1^*), \dots, P_k(\mathbf{w}_k^*)) \quad (10)$$

is a proximal map (from \mathbb{R}^{dk} to \mathbb{R}^{dk}).

Proof. Define $P(\mathbf{w}) = \sum_i \alpha_i P_i(\mathbf{w})$ if each $P_i(\mathbf{w})$ is single-valued. It is known that any proximal map P_i is almost everywhere single-valued [38], thus P is almost everywhere defined. Then, if necessary we take the closure of the graph of P so that it is defined everywhere. (This last step can be omitted if we always take the upper or lower limits at any jump of P_i .) With this interpretation of the average in (9), the rest of the first claim then follows from [Proposition 4, 48].

For the second claim about the product map in (10), let P_i be the proximal map of $f_i(\mathbf{w}_i^*)$. Then, it follows immediately from the definition (8) that the product map P is the proximal map of the sum function $f(\mathbf{w}^*) := \sum_i f_i(\mathbf{w}_i^*)$, where $\mathbf{w}^* = (\mathbf{w}_1^*, \dots, \mathbf{w}_k^*)$. \square

A.2 Proofs for Section 4

Let us first record a general result for iterates of the following type:

$$\mathbf{w}_{t+1}^* = \mathbf{w}_t^* + \eta_t \mathbf{z}_t^*, \quad (39)$$

where \mathbf{w}_0^* is given and \mathbf{z}_t^* may be arbitrary.

Lemma A.1. For any \mathbf{w} , any sequence of \mathbf{w}_t , and arbitrary function g , the iterates (39) satisfy:

$$\sum_{\tau=s}^t \eta_\tau [\langle \mathbf{w}_\tau - \mathbf{w}, -\mathbf{z}_\tau^* \rangle + g(\mathbf{w}_\tau) - g(\mathbf{w})] = \delta_{s-1}(\mathbf{w}) - \delta_t(\mathbf{w}) + \sum_{\tau=s}^t \delta_\tau(\mathbf{w}_\tau), \quad (40)$$

where $\delta_\tau(\mathbf{w}) := \frac{1}{\pi_\tau} g(\mathbf{w}) - \frac{1}{\pi_\tau} g(\mathbf{w}_{\tau+1}) - \langle \mathbf{w} - \mathbf{w}_{\tau+1}, \mathbf{w}_{\tau+1}^* \rangle$ and $\frac{1}{\pi_t} = 1 + \sum_{\tau=1}^t \eta_\tau$.

Proof. The proof is simple algebra (and was discovered by abstracting the original, tedious proof of Theorem 4.1). For any \mathbf{w} , we verify:

$$\sum_{\tau=s}^t \eta_\tau [\langle \mathbf{w}_\tau - \mathbf{w}, -\mathbf{z}_\tau^* \rangle + g(\mathbf{w}_\tau) - g(\mathbf{w})] \quad (41)$$

$$= \sum_{\tau=s}^t \langle \mathbf{w}_\tau - \mathbf{w}, \mathbf{w}_\tau^* - \mathbf{w}_{\tau+1}^* \rangle + \eta_\tau [g(\mathbf{w}_\tau) - g(\mathbf{w})] \quad (42)$$

$$= \sum_{\tau=s}^t \langle \mathbf{w}_\tau - \mathbf{w}, \mathbf{w}_\tau^* - \mathbf{w}_{\tau+1}^* \rangle + \left(\frac{1}{\pi_\tau} - \frac{1}{\pi_{\tau-1}} \right) [g(\mathbf{w}_\tau) - g(\mathbf{w})] \quad (43)$$

$$= \sum_{\tau=s}^t - \langle \mathbf{w}_\tau - \mathbf{w}_{\tau+1}, \mathbf{w}_{\tau+1}^* \rangle + \langle \mathbf{w} - \mathbf{w}_{\tau+1}, \mathbf{w}_{\tau+1}^* \rangle - \langle \mathbf{w} - \mathbf{w}_\tau, \mathbf{w}_\tau^* \rangle + \left(\frac{1}{\pi_\tau} - \frac{1}{\pi_{\tau-1}} \right) [g(\mathbf{w}_\tau) - g(\mathbf{w})] \quad (44)$$

$$= \sum_{\tau=s}^t \delta_\tau(\mathbf{w}_\tau) - \delta_\tau(\mathbf{w}) + \delta_{\tau-1}(\mathbf{w}). \quad (45)$$

Telescoping completes the proof. \square

Theorem 4.1. Suppose r^* is L -smooth (i.e. ∇r^* is L -Lipschitz continuous), then for any \mathbf{w} :

$$\sum_{\tau=0}^t \frac{\lambda_\tau}{\pi_\tau} [(\ell^{**} + r^{**})(\mathbf{w}_\tau) - (\ell^{**} + r^{**})(\mathbf{w})] \leq (1 - \lambda_0)\Delta(\mathbf{w}, \mathbf{w}_0) + \sum_{\tau=0}^t \frac{\lambda_\tau^2}{2\pi_\tau} L \|\mathbf{w}_\tau^* - \mathbf{z}_\tau^*\|_2^2, \quad (20)$$

where $\mathbf{w}_t := \nabla r^*(\mathbf{w}_t^*)$, $\pi_t := \prod_{\tau=1}^t (1 - \lambda_\tau)$, and $\pi_0 := 1$. $\Delta(\mathbf{w}, \mathbf{w}_t) := r^{**}(\mathbf{w}) - r^{**}(\mathbf{w}_t) - \langle \mathbf{w} - \mathbf{w}_t, \mathbf{w}_t^* \rangle$ is the Bregman divergence induced by the convex function r^{**} .

Proof. We first expand the recursion

$$\mathbf{w}_{t+1}^* = (1 - \lambda_t)\mathbf{w}_t^* + \lambda_t \mathbf{z}_t^* \quad (46)$$

$$= (1 - \lambda_t)(1 - \lambda_{t-1})\mathbf{w}_{t-1}^* + \lambda_t \mathbf{z}_t^* + (1 - \lambda_t)\lambda_{t-1}\mathbf{z}_{t-1}^* \quad (47)$$

$$= \dots \quad (48)$$

$$= (1 - \lambda_0)\pi_t \mathbf{w}_0^* + \sum_{\tau=0}^t \frac{\pi_t}{\pi_\tau} \lambda_\tau \mathbf{z}_\tau^*, \quad (49)$$

where $\pi_t := \prod_{\tau=1}^t (1 - \lambda_\tau)$ for $t \geq 1$ while $\pi_0 := 1$ and $\lambda_t \in [0, 1]$. We deduce (by, for instance, setting $\mathbf{z}_t^* = \mathbf{w}_t^* \equiv 1$, for all $t \geq 0$, in (49)) that

$$\frac{1}{\pi_t} = (1 - \lambda_0) + \sum_{\tau=0}^t \frac{\lambda_\tau}{\pi_\tau} = 1 + \sum_{\tau=1}^t \frac{\lambda_\tau}{\pi_\tau}. \quad (50)$$

Define $\frac{1}{\pi_{t-1}} := 1 - \lambda_0$ and $\mathbf{z}_t^* = -\nabla \ell^{**}(\mathbf{w}_t)$. We apply the convexity of ℓ^{**} to the LHS of (20):

$$\text{LHS} = \sum_{\tau=0}^t \frac{\lambda_\tau}{\pi_\tau} [\ell^{**}(\mathbf{w}_\tau) + r^{**}(\mathbf{w}_\tau) - \ell^{**}(\mathbf{w}) - r^{**}(\mathbf{w})] \quad (51)$$

$$\leq \sum_{\tau=0}^t \frac{\lambda_\tau}{\pi_\tau} [\langle \mathbf{w}_\tau - \mathbf{w}, -\mathbf{z}_\tau^* \rangle + r^{**}(\mathbf{w}_\tau) - r^{**}(\mathbf{w})]. \quad (52)$$

Next, we identify $\eta_\tau := \frac{\lambda_\tau}{\pi_\tau}$, $g = r^{**}$, $\mathbf{w}_{\tau+1}^* := \mathbf{w}_{\tau+1}^*/\pi_\tau = \mathbf{w}_\tau^* + \eta_\tau \mathbf{z}_\tau^*$, and $\delta_\tau = \frac{1}{\pi_\tau} \Delta$ so that we can apply Lemma A.1:

$$\text{LHS} \leq \sum_{\tau=0}^t \frac{\lambda_\tau}{\pi_\tau} [\langle \mathbf{w}_\tau - \mathbf{w}, -\mathbf{z}_\tau^* \rangle + r^{**}(\mathbf{w}_\tau) - r^{**}(\mathbf{w})] \quad (53)$$

$$\leq (1 - \lambda_0)\Delta(\mathbf{w}, \mathbf{w}_0) - \frac{1}{\pi_t}\Delta(\mathbf{w}, \mathbf{w}_{t+1}) + \sum_{\tau=0}^t \frac{1}{\pi_\tau}\Delta(\mathbf{w}_\tau, \mathbf{w}_{\tau+1}) \quad (54)$$

$$\leq (1 - \lambda_0)\Delta(\mathbf{w}, \mathbf{w}_0) + \sum_{\tau=0}^t \frac{1}{\pi_\tau} \frac{L}{2} \|\mathbf{w}_{\tau+1}^* - \mathbf{w}_\tau^*\|_2^2, \quad (55)$$

where in the last step we applied the nonnegativity of the Bregman divergence Δ (when induced by a convex function such as r^{**}) as well as the following inequality:

$$\Delta(\mathbf{w}_\tau, \mathbf{w}_{\tau+1}) = r^{**}(\mathbf{w}_\tau) - r^{**}(\mathbf{w}_{\tau+1}) - \langle \mathbf{w}_\tau - \mathbf{w}_{\tau+1}, \mathbf{w}_{\tau+1}^* \rangle \quad (56)$$

$$= r^{**}(\mathbf{w}_\tau) - r^{**}(\mathbf{w}_{\tau+1}) - \langle \mathbf{w}_\tau - \mathbf{w}_{\tau+1}, \nabla r^{**}(\mathbf{w}_{\tau+1}) \rangle \quad (\mathbf{w}_t = \nabla r^*(\mathbf{w}_t^*) \iff \mathbf{w}_t^* = \nabla r^{**}(\mathbf{w}_t)) \quad (57)$$

$$= r^{***}(\mathbf{w}_{\tau+1}^*) - r^{***}(\mathbf{w}_\tau^*) - \langle \mathbf{w}_{\tau+1}^* - \mathbf{w}_\tau^*, \nabla r^{***}(\nabla r^{**}(\mathbf{w}_\tau)) \rangle \quad (\text{by duality of Bregman divergence}) \quad (58)$$

$$= r^*(\mathbf{w}_{\tau+1}^*) - r^*(\mathbf{w}_\tau^*) - \langle \mathbf{w}_{\tau+1}^* - \mathbf{w}_\tau^*, \nabla r^*(\nabla r^{**}(\mathbf{w}_\tau)) \rangle \quad (\text{by convexity of } r^{***}) \quad (59)$$

$$= r^*(\mathbf{w}_{\tau+1}^*) - r^*(\mathbf{w}_\tau^*) - \langle \mathbf{w}_{\tau+1}^* - \mathbf{w}_\tau^*, \mathbf{w}_\tau \rangle \quad (\text{since } \nabla r^{**} = (\nabla r^*)^{-1}) \quad (60)$$

$$\leq \langle \mathbf{w}_{\tau+1}^* - \mathbf{w}_\tau^*, \nabla r^*(\mathbf{w}_\tau^*) \rangle + \frac{L}{2} \|\mathbf{w}_{\tau+1}^* - \mathbf{w}_\tau^*\|_2^2 - \langle \mathbf{w}_{\tau+1}^* - \mathbf{w}_\tau^*, \mathbf{w}_\tau \rangle \quad (\text{by smoothness of } r^*) \quad (61)$$

$$= \frac{L}{2} \|\mathbf{w}_{\tau+1}^* - \mathbf{w}_\tau^*\|_2^2 \quad (\text{since } \mathbf{w}_\tau := \nabla r^*(\mathbf{w}_\tau^*)). \quad (62)$$

Applying (46) completes the proof. \square

Corollary 4.2. Let $\bar{\mathbf{w}}_t := \sum_{\tau=0}^t \Lambda_{t,\tau} \mathbf{w}_\tau$, where $\Lambda_{t,\tau} := \frac{\lambda_\tau}{\pi_\tau} / H_t$ and $H_t := \sum_{\tau=0}^t \frac{\lambda_\tau}{\pi_\tau}$. Then, we have for any \mathbf{w} :

$$(\ell^{**} + r^{**})(\bar{\mathbf{w}}_t) - (\ell^{**} + r^{**})(\mathbf{w}) \leq \frac{(1 - \lambda_0)\Delta(\mathbf{w}, \mathbf{w}_0)}{H_t} + \frac{L}{2} \sum_{\tau=0}^t \lambda_\tau \Lambda_{t,\tau} \|\mathbf{w}_\tau^* - \mathbf{z}_\tau^*\|_2^2. \quad (21)$$

Proof. By the convexity of r^{**} and ℓ^{**} as well as the fact that the sum of two convex functions is convex, we have:

$$(\ell^{**} + r^{**})(\bar{\mathbf{w}}_t) \leq \sum_{\tau=0}^t \Lambda_{t,\tau} (\ell^{**} + r^{**})(\mathbf{w}_\tau). \quad (63)$$

Inserting the above in Theorem 4.1 (multiplied $1/H_t$) and noting that $\sum_{\tau=0}^t \Lambda_{t,\tau} = 1$, we have

$$(\ell^{**} + r^{**})(\bar{\mathbf{w}}_t) - (\ell^{**} + r^{**})(\mathbf{w}) \leq \sum_{\tau=0}^t \Lambda_{t,\tau} [(\ell^{**} + r^{**})(\mathbf{w}_\tau) - (\ell^{**} + r^{**})(\mathbf{w})] \quad (64)$$

$$\leq \frac{(1-\lambda_0)\Delta(\mathbf{w}, \mathbf{w}_0)}{H_t} + \frac{L}{2} \sum_{\tau=0}^t \lambda_\tau \Lambda_{t,\tau} \|\mathbf{w}_\tau^* - \mathbf{z}_\tau^*\|_2^2. \quad (65)$$

□

The following instantiation of Corollary 4.2 is notable: setting $\lambda_t = \frac{1}{t+1}$ implies $\pi_t = \lambda_t$, $H_t = t+1$, and therefore $\bar{\mathbf{w}}_t$ is simply ergodic averaging, i.e.,

$$\bar{\mathbf{w}}_t = \frac{1}{t+1} \sum_{s=0}^t \mathbf{w}_s. \quad (66)$$

The right-hand side of (21) diminishes at the rate of $O(\frac{\log t}{t})$ since

$$\sum_{\tau=0}^t \lambda_\tau \Lambda_{t,\tau} = \frac{1}{t+1} \sum_{\tau=0}^t \frac{1}{\tau+1}, \quad \text{and} \quad \sum_{\tau=0}^t \frac{1}{\tau+1} = O(\log t). \quad (67)$$

We remark that the log factor can be removed if we set $\lambda_t = \frac{2}{t+2}$ instead, for which we have

$$\pi_t = \prod_{\tau=1}^t \frac{\tau}{\tau+2} = \frac{2}{(t+1)(t+2)}, \quad (68)$$

$$H_t = \sum_{\tau=0}^t \frac{\lambda_\tau}{\pi_\tau} = \sum_{\tau=0}^t \frac{2}{\tau+2} \frac{(\tau+1)(\tau+2)}{2} = \frac{1}{2}(t+1)(t+2), \quad (69)$$

$$\sum_{\tau=0}^t \lambda_\tau \Lambda_{t,\tau} = \frac{4}{(t+1)(t+2)} \sum_{\tau=0}^t \frac{\tau+1}{\tau+2}, \quad \text{and} \quad \sum_{\tau=0}^t \frac{\tau+1}{\tau+2} = \mathcal{O}(t). \quad (70)$$

Proposition A.2.

$$\nabla \mathcal{M}_{r^*}^\mu(\mathbf{w}_t^*) = \mathbf{P}_{r^{**}}^{1/\mu}(\mathbf{w}_t^*/\mu) \quad (71)$$

Proof. By the envelope theorem, we have $\nabla \mathcal{M}_{r^*}^\mu(\mathbf{w}_t^*) = \frac{\mathbf{w}_t^* - \mathbf{P}_{r^*}^\mu(\mathbf{w}_t^*)}{\mu}$. Furthermore, using the Moreau decomposition, we have

$$\mathbf{w}_t^* = \mathbf{P}_{\mu r^*}^1(\mathbf{w}_t^*) + \mathbf{P}_{(\mu r^*)^*}^1(\mathbf{w}_t^*) \quad (72)$$

$$= \mathbf{P}_{r^*}^\mu(\mathbf{w}_t^*) + \mathbf{P}_s^\mu(\mathbf{w}_t^*), \quad (73)$$

where $s(\mathbf{w}) := r^{**}(\mathbf{w}/\mu)$. Combining the two results, we have

$$\nabla \mathcal{M}_{r^*}^\mu(\mathbf{w}_t^*) = \frac{\mathbf{w}_t^* - \mathbf{P}_{r^*}^\mu(\mathbf{w}_t^*)}{\mu} \quad (74)$$

$$= \frac{\mathbf{P}_s^\mu(\mathbf{w}_t^*)}{\mu} \quad (75)$$

$$= \frac{1}{\mu} \arg \min_{\mathbf{w}} \frac{1}{2\mu} \|\mathbf{w} - \mathbf{w}_t^*\|_2^2 + s(\mathbf{w}) \quad (76)$$

$$= \arg \min_{\mathbf{w}} \frac{1}{2\mu} \|\mu\mathbf{w} - \mathbf{w}_t^*\|_2^2 + s(\mu\mathbf{w}) \quad (77)$$

$$= \arg \min_{\mathbf{w}} \frac{\mu}{2} \|\mathbf{w} - \frac{\mathbf{w}_t^*}{\mu}\|_2^2 + r^{**}(\mathbf{w}) \quad (78)$$

$$= \mathbf{P}_{r^{**}}^{1/\mu}(\mathbf{w}_t^*/\mu). \quad (79)$$

□

A.3 Proofs for Section 5

In fact, Theorem 5.1 is a special case of a more general result:

Theorem A.3. Given \mathbf{w}_0^* , $\eta_t \geq 0$ and $\frac{1}{\pi_t} = 1 + \sum_{\tau=1}^t \eta_\tau$, consider the iterates defined as

$$\mathbf{w}_t = \mathbf{P}_r^{1/\mu_t}(\pi_{t-1}\mathbf{w}_t^*/\mu_t), \quad \mathbf{w}_{t+1}^* = \mathbf{w}_t^* - \eta_t \widetilde{\nabla} \ell(\mathbf{w}_t). \quad (80)$$

Then, for any \mathbf{w} ,

$$\sum_{\tau=s}^t \eta_\tau [\langle \mathbf{w}_\tau - \mathbf{w}, \widetilde{\nabla} \ell(\mathbf{w}_\tau) \rangle + r(\mathbf{w}_\tau) - r(\mathbf{w})] \leq \Delta_{s-1}(\mathbf{w}) - \Delta_t(\mathbf{w}) + \sum_{\tau=s}^t \Delta_\tau(\mathbf{w}_\tau) + \quad (81)$$

$$+ \sum_{\tau=s}^t \frac{1}{2} \left(\frac{\mu_{\tau+1}}{\pi_\tau} - \frac{\mu_\tau}{\pi_{\tau-1}} \right) (\|\mathbf{w}\|_2^2 - \|\mathbf{w}_\tau\|_2^2), \quad (82)$$

where $\Delta_\tau(\mathbf{w}) := r_\tau(\mathbf{w}) - r_\tau(\mathbf{w}_{\tau+1}) - \langle \mathbf{w} - \mathbf{w}_{\tau+1}, \mathbf{w}_{\tau+1}^* \rangle$ is the Bregman divergence induced by the (possibly nonconvex) function $r_\tau(\mathbf{w}) := \frac{1}{\pi_\tau} r(\mathbf{w}) + \frac{\mu_{\tau+1}}{2\pi_\tau} \|\mathbf{w}\|_2^2$.

Proof. We simply apply Lemma A.1 with $\mathbf{z}_\tau^* := -\widetilde{\nabla} \ell(\mathbf{w}_\tau)$, $\mathbf{w}_\tau^* = \mathbf{w}_\tau^*$ and $g := r$, and note that

$$\delta_\tau(\mathbf{w}) = \Delta_\tau(\mathbf{w}) - \frac{\mu_{\tau+1}}{2\pi_\tau} (\|\mathbf{w}\|_2^2 - \|\mathbf{w}_{\tau+1}\|_2^2). \quad (83)$$

To see that Δ_τ is the Bregman divergence of r_τ , we apply the (sub)differential optimality condition to

$$\mathbf{w}_{\tau+1} := \mathbf{P}_r^{1/\mu_{\tau+1}}(\pi_\tau \mathbf{w}_{\tau+1}^*/\mu_{\tau+1}) = \arg \min_{\mathbf{w}} \frac{1}{2} \|\mathbf{w} - \pi_\tau \mathbf{w}_{\tau+1}^*/\mu_{\tau+1}\|_2^2 + \frac{1}{\mu_{\tau+1}} r(\mathbf{w}), \quad (84)$$

so that

$$\mathbf{w}_{\tau+1}^* \in \frac{\mu_{\tau+1}}{\pi_\tau} \mathbf{w}_{\tau+1} + \frac{1}{\pi_\tau} \nabla r(\mathbf{w}_{\tau+1}) = \nabla r_\tau(\mathbf{w}_{\tau+1}) \quad (85)$$

and hence

$$\Delta_\tau(\mathbf{w}) = r_\tau(\mathbf{w}) - r_\tau(\mathbf{w}_{\tau+1}) - \langle \mathbf{w} - \mathbf{w}_{\tau+1}, \nabla r_\tau(\mathbf{w}_{\tau+1}) \rangle. \quad (86)$$

Clearly, r_τ is convex if r is convex, in which case $\Delta_\tau \geq 0$. □

Theorem 5.1. Fix any \mathbf{w} , the iterates in (31) satisfy:

$$\sum_{\tau=s}^t \eta_\tau [\langle \mathbf{w}_\tau - \mathbf{w}, \widetilde{\nabla} \ell(\mathbf{w}_\tau) \rangle + r(\mathbf{w}_\tau) - r(\mathbf{w})] \leq \Delta_{s-1}(\mathbf{w}) - \Delta_t(\mathbf{w}) + \sum_{\tau=s}^t \Delta_\tau(\mathbf{w}_\tau), \quad (32)$$

where $\Delta_\tau(\mathbf{w}) := r_\tau(\mathbf{w}) - r_\tau(\mathbf{w}_{\tau+1}) - \langle \mathbf{w} - \mathbf{w}_{\tau+1}, \mathbf{w}_{\tau+1}^* \rangle$ is the Bregman divergence induced by the (possibly nonconvex) function $r_\tau(\mathbf{w}) := \frac{1}{\pi_\tau} r(\mathbf{w}) + \frac{1}{2} \|\mathbf{w}\|_2^2$.

Proof. Simply set $\mu_t = \pi_{t-1}$ in Theorem A.3 above. \square

Corollary 5.2. For convex ℓ and any \mathbf{w} , the iterates in (31) satisfy:

$$\min_{\tau=s, \dots, t} \mathbf{E}[f(\mathbf{w}_\tau) - f(\mathbf{w})] \leq \frac{1}{\sum_{\tau=s}^t \eta_\tau} \cdot \mathbf{E}[\Delta_{s-1}(\mathbf{w}) - \Delta_t(\mathbf{w}) + \sum_{\tau=s}^t \Delta_\tau(\mathbf{w}_\tau)]. \quad (33)$$

If r is also convex, then

$$\min_{\tau=s, \dots, t} \mathbf{E}[f(\mathbf{w}_\tau) - f(\mathbf{w})] \leq \frac{1}{\sum_{\tau=s}^t \eta_\tau} \cdot \mathbf{E}[\Delta_{s-1}(\mathbf{w}) + \sum_{\tau=s}^t \frac{\eta_\tau^2}{2} \|\tilde{\nabla} \ell(\mathbf{w}_\tau)\|_2^2], \quad (34)$$

and

$$\mathbf{E}[f(\bar{\mathbf{w}}_t) - f(\mathbf{w})] \leq \frac{1}{\sum_{\tau=s}^t \eta_\tau} \cdot \mathbf{E}[\Delta_{s-1}(\mathbf{w}) + \sum_{\tau=s}^t \frac{\eta_\tau^2}{2} \|\tilde{\nabla} \ell(\mathbf{w}_\tau)\|_2^2], \quad (35)$$

where $\mathbf{w}_t = \frac{\sum_{\tau=s}^t \eta_\tau \mathbf{w}_\tau}{\sum_{\tau=s}^t \eta_\tau}$.

Proof. We first apply the expectation with respect to random sampling to (32) which reduces the left hand side to

$$\mathbf{E} \left[\sum_{\tau=s}^t \eta_\tau [\langle \mathbf{w}_\tau - \mathbf{w}, \tilde{\nabla} \ell(\mathbf{w}_\tau) \rangle + r(\mathbf{w}_\tau) - r(\mathbf{w})] \right] = \mathbf{E} \left[\sum_{\tau=s}^t \eta_\tau [\langle \mathbf{w}_\tau - \mathbf{w}, \nabla \ell(\mathbf{w}_\tau) \rangle + r(\mathbf{w}_\tau) - r(\mathbf{w})] \right] \quad (87)$$

$$\geq \mathbf{E} \left[\sum_{\tau=s}^t \eta_\tau [\ell(\mathbf{w}_\tau) - \ell(\mathbf{w}) + r(\mathbf{w}_\tau) - r(\mathbf{w})] \right] \quad (88)$$

$$= \mathbf{E} \left[\sum_{\tau=s}^t \eta_\tau [f(\mathbf{w}_\tau) - f(\mathbf{w})] \right], \quad (89)$$

where we used the convexity of ℓ . We then obtain (33) by using

$$\min_{\tau=s, \dots, t} \mathbf{E}[f(\mathbf{w}_\tau) - f(\mathbf{w})] \leq \frac{1}{\sum_{\tau=s}^t \eta_\tau} \mathbf{E} \left[\sum_{\tau=s}^t \eta_\tau [f(\mathbf{w}_\tau) - f(\mathbf{w})] \right]. \quad (90)$$

The right-hand sides of (34) and (35) are obtained by setting $\mu_\tau = \pi_{\tau-1}$ and upper bounding the Bregman divergence:

$$\Delta_\tau(\mathbf{w}_\tau) = r_\tau(\mathbf{w}_\tau) - r_\tau(\mathbf{w}_{\tau+1}) - \langle \mathbf{w}_\tau - \mathbf{w}_{\tau+1}, \mathbf{w}_{\tau+1}^* \rangle \quad (91)$$

$$= r_\tau^*(\mathbf{w}_{\tau+1}^*) - r_\tau^*(\mathbf{w}_\tau^*) - \langle \mathbf{w}_{\tau+1}^* - \mathbf{w}_\tau^*, \mathbf{w}_\tau \rangle \quad (\text{by duality of Bregman divergence}) \quad (92)$$

$$\leq \frac{1}{2} \|\mathbf{w}_\tau^* - \mathbf{w}_{\tau+1}^*\|_2^2 \quad (\text{by 1-smoothness of } r^*) \quad (93)$$

$$= \frac{\eta_t^2}{2} \|\tilde{\nabla} \ell(\mathbf{w}_t)\|_2^2, \quad (94)$$

where we have used the well-known fact that the Fenchel conjugate of a $(1/L)$ -strongly convex function is L -smooth (in our case $L = 1$). Lastly, the left-hand side of (35) is obtained by applying the convexity of f :

$$\left[\sum_{\tau=s}^t \eta_\tau \right] f(\bar{\mathbf{w}}_t) \leq \sum_{\tau=s}^t \eta_\tau f(\mathbf{w}_\tau), \quad \text{where } \bar{\mathbf{w}}_t = \frac{\sum_{\tau=s}^t \eta_\tau \mathbf{w}_\tau}{\sum_{\tau=s}^t \eta_\tau}. \quad (95)$$

\square

A.3.1 Discussion of Theorem A.3

Recall the Bregman divergence from Theorem A.3:

$$\Delta_\tau(\mathbf{w}) = r_\tau(\mathbf{w}) - r_\tau(\mathbf{w}_{\tau+1}) - \langle \mathbf{w} - \mathbf{w}_{\tau+1}, \mathbf{w}_{\tau+1}^* \rangle, \quad r_\tau(\mathbf{w}) = \frac{1}{\pi_\tau} r(\mathbf{w}) + \frac{\mu_{\tau+1}}{2\pi_\tau} \|\mathbf{w}\|_2^2. \quad (96)$$

When r is σ_0 -strongly convex, r_τ is $\frac{\sigma_0 + \mu_{\tau+1}}{\pi_\tau}$ -strongly convex, and hence

$$\Delta_\tau(\mathbf{w}_\tau) \leq \frac{\pi_\tau}{2(\sigma_0 + \mu_{\tau+1})} \|\mathbf{w}_{\tau+1}^* - \mathbf{w}_\tau^*\|_2^2 = \frac{\pi_\tau \eta_\tau^2}{2(\sigma_0 + \mu_{\tau+1})} \|\tilde{\nabla} \ell(\mathbf{w}_\tau)\|_2^2, \quad (97)$$

where we used the duality of the Bregman divergence and the smoothness of r_τ^* . Dividing both sides of (81) by $\sum_{\tau=s}^t \eta_\tau$ we obtain the upper bound:

$$\text{UB} := \frac{\frac{\pi_{s-1}}{\sigma_0 + \mu_s} \|\mathbf{w}^* - \mathbf{w}_s^*\|_2^2 + \sum_{\tau=s}^t \left[\frac{\pi_\tau \eta_\tau^2}{\sigma_0 + \mu_{\tau+1}} \|\tilde{\nabla} \ell(\mathbf{w}_\tau)\|_2^2 + \left(\frac{\mu_{\tau+1}}{\pi_\tau} - \frac{\mu_\tau}{\pi_{\tau-1}} \right) (\|\mathbf{w}\|_2^2 - \|\mathbf{w}_\tau\|_2^2) \right]}{2 \sum_{\tau=s}^t \eta_\tau}, \quad (98)$$

where $\mathbf{w}^* := \nabla r_{s-1}(\mathbf{w})$ and we have dropped the non-positive term $-\Delta_t(\mathbf{w})$. Suppose⁸ $\frac{\mu_{t+1}}{\pi_t}$ is non-decreasing w.r.t. t , we can thus drop some non-positive terms to further simplify:

$$\text{UB} \leq \frac{\frac{\pi_{s-1}}{\sigma_0 + \mu_s} \|\mathbf{w}^* - \mathbf{w}_s^*\|_2^2 + \frac{\mu_{t+1}}{\pi_t} \|\mathbf{w}\|_2^2 + \sum_{\tau=s}^t \left[\frac{\pi_\tau \eta_\tau^2}{\sigma_0 + \mu_{\tau+1}} \|\tilde{\nabla} \ell(\mathbf{w}_\tau)\|_2^2 \right]}{2 \sum_{\tau=s}^t \eta_\tau}, \quad (99)$$

To minimize the upper bound, we consider two cases:

- $\sigma_0 = 0$, in which case let us choose $\mu_{\tau+1} = c\sqrt{\eta_\tau \pi_\tau} = c\sqrt{\lambda_\tau}$ (recall that we reparameterized $\eta_\tau = \lambda_\tau / \pi_\tau$ from GCG), where c is an absolute constant. Then, the upper bound reduces to

$$\frac{\frac{\pi_{s-1}}{\mu_s} \|\mathbf{w}^* - \mathbf{w}_s^*\|_2^2 + \frac{c\eta_t}{\sqrt{\lambda_t}} \|\mathbf{w}\|_2^2 + \sum_{\tau=s}^t \eta_\tau \sqrt{\lambda_\tau} \|\tilde{\nabla} \ell(\mathbf{w}_\tau)\|_2^2 / c}{2 \sum_{\tau=s}^t \eta_\tau}, \quad (100)$$

where recall that $\pi_t = \frac{1}{1 + \sum_{\tau=1}^t \eta_\tau}$. When the gradient $\nabla \ell$ is bounded, we may choose

$$\lambda_t = \frac{\eta_t}{1 + \sum_{\tau=1}^t \eta_\tau} = O(1/t) \text{ and } \frac{1}{\sum_{\tau=1}^t \eta_\tau} = O(1/\sqrt{t}) \quad (101)$$

so that the upper bound diminishes⁹ at the rate of $O(1/\sqrt{t})$. This result makes intuitive sense, since we know the optimal step size λ_t in GCG is $\Theta(1/t)$. Also, it reveals that the smoothing parameter $\mu_t = O(1/\sqrt{t})$, matching the rate of the upper bound (i.e., progress on the original problem). Interestingly, the choice in (101) can be realized in multiple ways. In fact, $\eta_t = \eta_0 t^{-p}$ for any $p \in [0, \frac{1}{2}]$ suffices. In particular, choosing $p = 0$ leads to the constant step size $\eta_t \equiv \eta_0$. However, note that these choices are in some sense equivalent, since they all lead to $\lambda_t = O(1/t)$ and $\mu_t = O(1/\sqrt{t})$ hence the underlying GCG progresses similarly.

- $\sigma_0 > 0$, in which case we can further relax the upper bound to:

$$\text{UB} \leq \frac{\frac{\pi_{s-1}}{\sigma_0 + \mu_{s-1}} \|\mathbf{w}^* - \mathbf{w}_s^*\|_2^2 + \frac{\mu_{t+1}}{\pi_t} \|\mathbf{w}\|_2^2 + \sum_{\tau=s}^t \left[\frac{\pi_\tau \eta_\tau^2}{\sigma_0} \|\tilde{\nabla} \ell(\mathbf{w}_\tau)\|_2^2 \right]}{2 \sum_{\tau=s}^t \eta_\tau}, \quad (102)$$

and now we may choose $\mu_{\tau+1} = c\eta_\tau \pi_\tau = c\lambda_\tau$, which approaches 0 significantly faster than before. With this choice, the upper bound reduces to

$$\frac{\frac{\pi_{s-1}}{\sigma_0 + \mu_{s-1}} \|\mathbf{w}^* - \mathbf{w}_s^*\|_2^2 + c\eta_t \|\mathbf{w}\|_2^2 + \sum_{\tau=s}^t \eta_\tau \lambda_\tau [\|\tilde{\nabla} \ell(\mathbf{w}_\tau)\|_2^2 / \sigma_0]}{2 \sum_{\tau=s}^t \eta_\tau}. \quad (103)$$

Again, we may set λ_t as in (101), and we can choose constant $\eta_t \equiv \eta_0$. The major difference is that we may now decrease the smoothing parameter μ_t much more aggressively.

⁸This assumption can be easily satisfied. In fact, we can just set $\mu_{t+1} = \pi_t$ (as in the main paper), which would simplify the discussion quite a bit.

⁹Note that by choosing $s \propto t$, the averaged sequence $\frac{\sum_{\tau=s}^t \eta_\tau \sqrt{\lambda_\tau}}{\sum_{\tau=s}^t \eta_\tau} \leq \max_{\tau=s, \dots, t} \sqrt{\lambda_\tau}$ diminishes similarly in order as $\sqrt{\lambda_t}$.

B Implementation and Experiment Details

B.1 CIFAR-10

B.1.1 Model and Quantization Details

Similar to Bai et al. [5], we use ResNets for which we quantize all weights. BatchNormalization layers and activations are kept at full precision. The basic ResNet implementation is taken from https://github.com/akamaster/pytorch_resnet_cifar10/blob/master/resnet.py.

B.1.2 Data Augmentation

We follow the data augmentation strategy from Bai et al. [5]: padding by four pixels on each side, randomly cropping to 32-by-32 pixels, horizontally flipping with probability one half. Finally, the images are normalized by subtracting (0.4914, 0.4822, 0.4465) and subsequently dividing by (0.247, 0.243, 0.261).

B.1.3 Pretrained Full Precision Model Setup

We pretrain two single full precision models for 200 epochs using the standard optimization setup: SGD with 0.9 momentum and $1e-4$ weight decay. The initial learning rate is 0.1 which is multiplied by 0.1 at epoch 100 and 150. We use batch size 128.

B.1.4 Fine-Tuning Setup

All methods are initialized with the last checkpoint of the full precision models.

For BinaryConnect, we train with the recommended strategy from Courbariaux et al. [11]: Adam with learning rate 0.01 and multiplication of the learning rate by 0.1 at epoch 81 and 122. For ProxQuant we use Adam with fixed learning rate of 0.01 as in Bai et al. [5]. We did not perform hyperparameter search over the optimization setup for ProxConnect but rather just used the optimization setup from BinaryConnect as the methods are very similar. We use batch size 128.

Similar to Bai et al. [5], we perform a hard quantization at epoch 200: all weights are projected to their closet quantization points. As in Bai et al. [5], we then train BatchNormalization layers for another 100 epochs.

B.1.5 End-To-End Setup

For simplicity and to avoid an expensive hyperparameter search, we followed the above full precision optimization setup for all methods. Similar to the fine-tuning setup, we perform a hard quantization at epoch 200 and keep training BatchNormalization layers for another 100 epochs. We use batch size 128.

B.1.6 Compute and Resources

We run all CIFAR-10 experiments on our internal cluster with GeForce GTX 1080 Ti. We use one GPU per experiment. The total amount of GPU hours is summarized in Table 5.

Table 5: Compute for CIFAR-10 experiments measured in hours per single GeForce GTX 1080 Ti GPU.

Architecture	Run time	# of total experiments	Total run time
ResNet20	1.5	181	271.5
ResNet56	3	181	543

B.1.7 Additional Results

As mentioned in Section 6.1, we perform a small grid search over ρ_0 . After an initial exploration stage, we found good regions of ρ_0 for ProxQuant and ProxConnect. Since ProxQuant and

reverseProxConnect are quite similar, in particular in the small ρ regime, we simply used the ρ_0 from ProxQuant for reverseProxConnect. See below the results for all ρ_0 settings.

ProxConnect is reasonably stable with respect to the choice of ρ_0 for both fine-tuning and end-to-end training. ProxQuant and reverseProxConnect, on the other hand, are very sensitive to the choice of ρ_0 for end-to-end training. ProxConnect reduces to BinaryConnect for large ρ , and therefore it is in line with our experiments that ProxConnect should be stable with respect to ρ_0 , particularly choosing ρ_0 too large should not be an issue.

Table 6: Additional results for fine-tuning ProxQuant.

Model	Quantization	$\rho_0 = 5e-7$	$\rho_0 = 1e-6$	$\rho_0 = 2e-6$
ResNet20	Binary	89.68 (0.10)	89.94 (0.10)	89.34 (0.33)
	Ternary	91.03 (0.07)	91.46 (0.05)	91.11 (0.10)
	Quaternary	91.10 (0.06)	91.13 (0.18)	91.43 (0.17)
ResNet56	Binary	92.25 (0.08)	92.33 (0.06)	92.16 (0.14)
	Ternary	92.52 (0.90)	93.07 (0.02)	92.71 (0.16)
	Quaternary	92.49 (0.03)	92.82 (0.13)	92.80 (0.15)

Table 7: Additional results for fine-tuning reverseProxConnect.

Model	Quantization	$\rho_0 = 5e-7$	$\rho_0 = 1e-6$	$\rho_0 = 2e-6$
ResNet20	Binary	89.88 (0.23)	89.98 (0.17)	89.91 (0.24)
	Ternary	91.30 (0.05)	91.36 (0.03)	91.47 (0.15)
	Quaternary	91.05 (0.12)	91.43 (0.05)	91.42 (0.07)
ResNet56	Binary	92.18 (0.06)	92.35 (0.02)	92.47 (0.29)
	Ternary	92.84 (0.09)	92.85 (0.12)	92.84 (0.17)
	Quaternary	92.66 (0.20)	92.88 (0.17)	92.91 (0.22)

Table 8: Additional results for fine-tuning ProxConnect.

Model	Quantization	$\rho_0 = 5e-3$	$\rho_0 = 1e-2$	$\rho_0 = 2e-2$
ResNet20	Binary	89.63 (0.26)	90.29 (0.07)	90.31 (0.21)
	Ternary	91.31 (0.07)	91.37 (0.18)	91.13 (0.27)
	Quaternary	91.62 (0.21)	91.55 (0.10)	91.81 (0.14)
ResNet56	Binary	92.41 (0.13)	92.62 (0.09)	92.65 (0.16)
	Ternary	93.17 (0.04)	93.25 (0.12)	93.22 (0.06)
	Quaternary	93.41 (0.11)	93.42 (0.12)	93.28 (0.06)

Table 9: Additional results end-to-end training ProxQuant.

Model	Quantization	$\rho_0 = 1e-7$	$\rho_0 = 1e-6$	$\rho_0 = 1e-5$
ResNet20	Binary	81.59 (0.75)	81.49 (0.41)	71.90 (0.75)
	Ternary	28.22 (1.70)	41.08 (2.95)	47.98 (1.06)
	Quaternary	84.58 (0.15)	85.29 (0.07)	75.08 (0.16)
ResNet56	Binary	86.13 (1.71)	80.25 (0.51)	68.31 (2.21)
	Ternary	21.93 (2.43)	41.11 (2.12)	50.54 (3.01)
	Quaternary	87.81 (1.30)	83.57 (1.70)	72.58 (2.23)

Table 10: Additional results end-to-end training reverseProxConnect.

Model	Quantization	$\rho_0 = 1e-7$	$\rho_0 = 1e-6$	$\rho_0 = 1e-5$
ResNet20	Binary	81.82 (0.32)	80.84 (0.40)	72.10 (1.01)
	Ternary	26.49 (2.82)	40.78 (0.39)	47.17 (1.94)
	Quaternary	85.05 (0.27)	84.82 (0.32)	75.61 (0.54)
ResNet56	Binary	86.25 (1.50)	81.58 (0.92)	67.53 (2.74)
	Ternary	21.51 (1.12)	42.95 (1.57)	36.34 (18.68)
	Quaternary	87.30 (1.02)	84.72 (1.31)	73.36 (1.36)

Table 11: Additional results end-to-end training ProxConnect.

Model	Quantization	$\rho_0 = 5e-3$	$\rho_0 = 1e-2$	$\rho_0 = 2e-2$
ResNet20	Binary	89.72 (0.13)	89.92 (0.26)	89.65 (0.15)
	Ternary	84.09 (0.16)	83.54 (0.36)	82.84 (0.34)
	Quaternary	90.17 (0.14)	90.12 (0.33)	89.91 (0.09)
ResNet56	Binary	91.26 (0.59)	90.45 (0.83)	89.29 (0.45)
	Ternary	84.36 (0.75)	83.46 (1.14)	82.54 (1.38)
	Quaternary	91.00 (0.50)	90.76 (0.54)	91.70 (0.14)

B.2 ImageNet

B.3 Model and Quantization Details

We use ResNet18 as our model, for which we quantize all weights except for the first convolutional layer and the last fully-connected layer. Other components such as BatchNormalization layers, activations and biases are kept at full-precision. The ResNet implementation is borrowed from PyTorch’s torchvision package: <https://github.com/pytorch/vision/blob/882e11db8138236ce375ea0dc8a53fd91f715a90/torchvision/models/resnet.py>.

B.4 Data Augmentation

Common data augmentation strategy is followed from <https://github.com/pytorch/examples/blob/c002856901eaf9be112feb9b14a9d5c3e779da74/imagenet/main.py#L204-L231>. At training time, the images are randomly resized and cropped to 224-by-224 pixels, followed by a random horizontal flip. At inference time, images are first resized to 256-by-256 pixels, then center cropped to 224-by-224 pixels. Finally, the images are normalized by subtracting (0.485, 0.456, 0.406) and subsequently dividing by (0.229, 0.224, 0.224).

B.5 Pretrained Full Precision Model Setup

We use the ResNet18 checkpoint provided by torchvision as our pretrained model: <https://download.pytorch.org/models/resnet18-f37072fd.pth>.

B.6 Fine-Tuning Setup

All fine-tuned models are initialized from the pretrained model described in the previous section.

As advised by Alizadeh et al. [2], we use Adam as the optimizer to fine-tune the models. We use Adam with default parameters, initial learning rate $1e-4$ and batch size 256. The models are fine-tuned for 50 epochs. The learning rate is divided by 10 at epoch 15 and 30. We perform hard quantization at epoch 45 and train the remaining full-precision layers for the last 5 epochs, mostly to let BatchNormalization layers stabilize.

B.7 End-To-End Setup

For the end-to-end setup, the models are trained from scratch mimicking the training setup of the full-precision pretrained model. The quantized ResNet18 is trained for 90 epochs using SGD with a starting learning rate of 0.1, momentum of 0.9 and weight decay of $1e-4$. The learning rate is multiplied by 0.1 at epoch 30 and 60. We use batch size 256. We perform hard quantization at epoch 80 and train the remaining full-precision layers for the last 10 epochs.

B.8 Compute and Resources

We run all experiments on our internal cluster with Tesla V100. We use one GPU per experiment. The total amount of GPU hours is summarized in Table 12.

Table 12: Compute for ImageNet experiments measured in hours per single Tesla V100 GPU.

Approach	Run time	# of total experiments	Total run time
End-To-End	27	9	243
Fine-Tuning	12	9	108