

AdaSTE: An Adaptive Straight-Through Estimator to Train Binary Neural Networks

Huu Le Rasmus Kjær Høier Che-Tsung Lin Christopher Zach
Chalmers University of Technology, Gothenburg, Sweden
huul, hier, chetsung, zach@chalmers.se

Abstract

We propose a new algorithm for training deep neural networks (DNNs) with binary weights. In particular, we first cast the problem of training binary neural networks (BiNNs) as a bilevel optimization instance and subsequently construct flexible relaxations of this bilevel program. The resulting training method shares its algorithmic simplicity with several existing approaches to train BiNNs, in particular with the straight-through gradient estimator successfully employed in BinaryConnect and subsequent methods. In fact, our proposed method can be interpreted as an adaptive variant of the original straight-through estimator that conditionally (but not always) acts like a linear mapping in the backward pass of error propagation. Experimental results demonstrate that our new algorithm offers favorable performance compared to existing approaches.¹

1. Introduction

Deploying deep neural networks (DNNs) to computing hardware such as mobile and IoT devices with limited computational and storage resources is becoming increasingly relevant in practice, and hence training methods especially dedicated to quantized DNNs have emerged as important research topics in recent years [9]. In this work, we are particularly interested in the special case of DNNs with binary weights limited to $\{+1, -1\}$, since in this setting the computations at inference time largely reduce to sole additions and subtractions. Very abstractly, the task of learning in such binary weight neural networks (BiNNs) can be formulated as an optimization program with binary constraints on the network parameters, i.e.,

$$\begin{aligned} \min_w \ell(w) \quad & \text{s.t. } w \in \{-1, 1\}^d, & (1) \\ = \min_{w \in \{-1, 1\}^d} \mathbb{E}_{(x, y) \sim p_{\text{data}}} [\psi(f(x, w), y)], & (2) \end{aligned}$$

¹This work was partially supported by the Wallenberg AI, Autonomous Systems and Software Program (WASP) funded by the Knut and Alice Wallenberg Foundation.

where d is the dimensionality of the underlying parameters (i.e. all network weights), p_{data} is the training distribution and ψ is the training loss (such as the cross-entropy or squared Euclidean error loss). $f(x; w)$ is the prediction of the DNN with weights w for input x .

In practice, one needs to address problem settings where the parameter dimension d is very large (such as deep neural networks with many layers). However, addressing the binary constraints in the above program is a challenging task, which is due to the combinatorial and non-differentiable nature of the underlying optimization problem. In view of large training datasets, (stochastic) gradient-based methods to obtain minimizers of (1) are highly preferable. Various techniques have been proposed to address the above difficulties and convert (1) into a differentiable surrogate. The general approach is to introduce real-valued “latent” weights $\theta \in \mathbb{R}^d$, from which the effective weights $w = \text{sgn}(\theta)$ are generated via the sign function (or a differentiable surrogate thereof). One of the simplest and nevertheless highly successful algorithms to train BiNNs termed BinaryConnect [10] is based on straight-through estimators (STE), which ignore the sign mapping entirely when forming the gradient w.r.t. the latent weights θ (and therefore the update of θ is based on $\nabla_w \ell(w)$ instead of $\nabla_\theta \ell(\text{sgn}(\theta))$). Although this appears initially not justified, BinaryConnect works surprisingly well and is still a valid baseline method for comparison. More recently, the flexibility in choosing the distance-like mapping leveraged in the mirror descent method [30] (and in particular the entropy descent algorithm [7]) provides some justification of BinaryConnect-like methods [3] (see also Sec. 3.2).

In this work, we propose a new framework for training binary neural networks. In particular, we first formulate the training problem shown in (1) as a bilevel optimization task, which is subsequently relaxed using an optimal value reformulation. Further, we propose a novel scheme to calculate meaningful gradient surrogates in order to update the network parameters. The resulting method strongly resembles BinaryConnect but leverages an adaptive variant of the straight-through gradient estimator: the sign func-




Figure 1. Adaptive straight-through estimation illustrated when $s = \tanh$. ℓ' is the incoming back-propagated error signal. Left: $\theta \approx 0$. The finite difference slope $(\hat{w} - w^*)/\beta$ matches the derivative of \tanh very well. Middle: $\theta \ll 0 \wedge \ell' < 0$. A nearly vanishing derivative of \tanh is boosted and \tanh becomes “leaky.” Right: $\theta \ll 0 \wedge \ell' > 0$. No gradient “boosting” in this case. The case $\theta \gg 0$ is symmetrical.

tion is conditionally replaced by a suitable linear but data-dependent mapping. Fig. 1 illustrates the underlying principle for the tanh mapping: depending on the incoming error signal, vanishing gradients induced by tanh are conditionally replaced by non-vanishing finite-difference surrogates. We finally point out that our proposed method can be cast as a mirror descent method using a data-dependent and varying distance-like mapping.

2. Related Work

The practical motivation for exploring weight quantization is to reduce the computational costs of deploying (and in some cases training) neural networks. This can be particularly attractive in the case of edge computing and IoT devices [9]. Even when retaining floating point precision for activations z , using binarized weights matrices W means that the omnipresent product Wz reduces to cheaper additions and subtractions of floating point values.

Already in the early 1990s, [14, 43] trained BiNNs using fully local learning rules with layerwise targets computed via node perturbations. In order to avoid the limited scalability of node perturbations, [37] instead employed a differentiable surrogate of the sign function for gradient computation. Recently the use of differentiable surrogates in the backwards pass has been coined the Backward Pass Differentiable Approximation (BPDA) in the context of adversarial attacks [5]. However, the same principle is at the core of many network quantization approaches, most notably the STE for gradient estimation.

Recent approaches have mainly focused on variations of the STE. A set of real valued (latent) weights are binarized when computing the forward pass, but during the backwards pass the identity mapping is used as its differentiable surrogate (which essentially makes the STE a special case of BPDA). The computed gradients are then used to update the latent weights. The STE was presented by Hinton (and accredited to Krizhevsky) in a video lecture in 2012 [19]. Subsequently it was employed for training networks with binary activations in [8], and to train networks with binary weights (and floating point activations) in the BinaryConnect (BC) model [10]. BinaryConnect also used

heuristics such as clipping the latent weights and employing Batch Normalization [22] (including its use at the output layer) to improve the performance of STE based training. Further and recent analysis of the straight-through estimator is provided in [47], where its origin is traced back to early work on perceptrons Rosenblatt [35, 36]. The STE has also been applied to training fully binarized neural networks (e.g. [21]). Moreover, Rastegari et al. [34] employ the STE for training fully binarized as well as mixed precision networks, and achieve improved performance by introducing layer and channel-wise scaling factors. An interesting line of research [13, 42] has explored adapting the STE to learn parameters of the quantization mapping (such as quantization steps and bit-width subjected to a memory budget).

Subsequent approaches have focused on deriving similar but less heuristic learning algorithms for networks with binary weights. ProxQuant (PQ) [6], Proximal Mean-Field (PMF) [2], Mirror Descent (MD) [3] and Rotated Binary Neural Networks (RBNN) [26] formulate the task of training DNNs with binary weights as a constrained optimization problem and propose different conversion functions used for moving between real-valued latent weights and binarized weights. A common feature among these methods is that they belong to the class of homotopy methods by gradually annealing the conversion mapping. Qin et al [33] introduce a novel technique for minimizing the information loss (caused by binarization) in the forward pass, and also aims to address gradient error by employing a gradually annealed tanh function as a differentiable surrogate during the backwards pass along with a carefully chosen gradient clipping schedule. Similar to early research, [18] does not introduce latent real-valued weights, but rather updates the binary weights directly using a momentum based optimizer designed specifically for BiNNs. Several authors have approached the training of quantized neural networks via a variational approach [1, 27, 29, 40]. Among those, Bayes-BiNN [29] is particularly competitive: instead of optimizing over binary weights, the parameters of Bernoulli distributions are learned by employing both a Bayesian learning rule [24] and the Gumbel-softmax trick [23, 28] (therefore requiring an inverse temperature parameter to convert the concrete distribution to a Bernoulli one).

For additional surveys of weight quantization we refer to the review papers [16, 32] as well as section III of [11]. For a review of the efficacy of various ad-hoc techniques commonly employed for training BiNNs we refer to [4].

3. Background

After clarifying some mathematical notations we summarize the mirror descent method (and its use to train BiNNs) and the Prox-Quant approach in order to better establish similarities and differences with our proposed method later.

3.1. Notation

A constraint such as $w \in C$ is written as $\iota_C(w)$ in functional form. We use \odot to denote element-wise multiplication and \oslash for element-wise division. The derivative of a function ℓ at w is written as $\ell'(w)$. Many mappings will be piece-wise differentiable but continuous. Therefore, in those cases $\ell'(w)$ is a suitable element in the sub- or super-derivative. We use an arrow over some variable names (especially $\vec{\beta}$) to emphasize that this is a vector and not a scalar. For the same reason we use e.g. \vec{s} and sgn to indicate the vectorized form of a scalar mapping s (or sgn) that is applied element-wise.

3.2. Mirror Descent

In short, mirror descent [7, 30] successively generates new iterates by minimizing a regularized first-order surrogate of the target objective. The most common quadratic regularizer (which leads to the gradient descent method) is replaced by a more general Bregman divergence penalizing large deviations from the previous iterate. The main motivation is to accelerate convergence of first-order methods, but it can also yield very elegant methods such as the entropic descent algorithm, where the utilized Bregman divergence based on the (negated) Shannon entropy is identical to the KL divergence. Entropic descent is very natural when optimizing unknowns constrained to remain in the probability simplex Δ . The algorithm repeats updates of the form

$$w^{(t+1)} \leftarrow \arg \min_{w \in \Delta} w^\top \ell'(w^{(t)}) + \frac{1}{\eta} D_{KL}(w \| w^{(t)}) \quad (3)$$

with the associated first-order optimality condition

$$w_j^{(t+1)} \propto w_j^{(t)} e^{-\eta \ell'(w^{(t)})_j}. \quad (4)$$

Reparametrizing w as $w = \sigma(\theta)$, where σ is the soft-arg-max function, $\sigma(u)_j = e^{u_j} / \sum_{j'} e^{u_{j'}}$, yields

$$\theta^{(t+1)} \leftarrow \theta^{(t)} - \eta \ell'(w^{(t)}) = \theta^{(t)} - \eta \ell'(\sigma(\theta^{(t)})). \quad (5)$$

Interestingly, mirror descent modifies the chain rule by bypassing the inner derivative, since the update is based on

$\ell'(\sigma(\theta^{(t)}))$ and not on $\frac{d}{d\theta} \ell(\sigma(\theta^{(t)}))$ as in regular gradient descent. Hence, mirror descent is one way to justify the straight-through estimator. The entropic descent algorithm is leveraged in [3] to train networks with binary (and also generally quantized) weights. The soft-arg-max function σ is slowly modified towards a hard arg-max mapping in order to ultimately obtain strictly quantized weights.

3.3. ProxQuant

ProxQuant [6] is based on the observation that the straight-through gradient estimator is linked to proximal operators via the dual averaging method [44]. The proximal operator for a function ϕ is the solution of the following least-squares regularized optimization problem,

$$\text{prox}_{\lambda\phi}(\theta) = \arg \min_{\theta'} \lambda\phi(\theta') + \frac{1}{2} \|\theta' - \theta\|^2, \quad (6)$$

where $\lambda > 0$ controls the regularization strength. If ϕ is a convex and lower semi-continuous mapping, the minimizer of the r.h.s. is always unique and $\text{prox}_{\lambda\phi}$ is a proper function (and plays an crucial role in many convex optimization methods). ProxQuant uses a non-convex mappings for ϕ , which is far more uncommon for proximal steps than the convex case (see e.g. [41] for another example). In order to train DNNs with binary weights, ϕ is chosen as W-shaped function,

$$\phi(\theta) = \sum_{j=1}^d \min \{|\theta_j - 1|, |\theta_j + 1|\}. \quad (7)$$

ϕ has 2^d isolated global minima and is therefore not convex. Note that $\text{prox}_{\lambda\phi}(\theta)$ is uniquely defined as long as all elements in θ are non-zero. The network weights are updated according to

$$\theta^{(t+1)} \leftarrow \text{prox}_{\lambda^{(t)}\phi}(\theta^{(t)} - \eta \ell'(\theta^{(t)})), \quad (8)$$

and the regularization weight $\lambda^{(t)}$ is increased via an annealing schedule, which makes ProxQuant an instance of homotopy methods: strictly quantized weights are only obtained for a sufficiently large value of $\lambda^{(t)}$.

4. Adaptive Straight-Through Estimator

In this section, we propose a new approach to tackle the optimization problem given in (1). Reformulating and relaxing an underlying bilevel minimization problem was initially inspired by contrastive Hebbian learning for DNNs [38, 45, 49] and is at the core of the proposed method.

4.1. Bilevel Optimization Formulation

We start by rewriting the original problem (1) as the following bilevel minimization program,

$$\min_{\theta} \ell(w^*) \quad \text{s.t. } w^* = \arg \min_w \mathcal{E}(w; \theta) \quad (9)$$

where $\mathcal{E}(w; \theta)$ can be any function that favors w^* to be binary. Two classical choices for \mathcal{E} are given by

$$\mathcal{E}_{\text{tanh}}(w; \theta) = -\frac{1}{\tau} \sum_j H\left(\frac{1}{2}(1 - w_j)\right) - w^\top \theta \quad (10)$$

$$\mathcal{E}_{\text{hard-tanh}}(w; \theta) = \frac{1}{2\tau} \|w\|^2 - w^\top \theta + \nu_{[-1,1]^d}(w), \quad (11)$$

where H is the Shannon entropy of a Bernoulli random variable, $H(u) = -u \log u - (1 - u) \log(1 - u)$. The minimizer w^* for given θ is the tanh mapping in the case of $\mathcal{E}_{\text{tanh}}$, $w_j^* = \tanh(\theta_j/\tau)$, and the second option yields the hard-tanh mapping, $w_j^* = \Pi_{[-1,1]}(\theta_j/\tau)$. $\tau > 0$ is a parameter steering how well these mappings approximate the sign function $\text{s\grave{g}n}(\theta)$.

In order to apply a gradient-based learning method we require that \mathcal{E} is differentiable w.r.t. θ for all w . In the above examples we have $\frac{\partial}{\partial \theta} \mathcal{E}(w; \theta) = -w$. It will be sufficient for our purposes to assume that \mathcal{E} is of the form

$$\mathcal{E}(w; \theta) = -w^\top \theta + \mathcal{G}(w) \quad (12)$$

for a coercive function \mathcal{G} bounded from below. That is, w and θ only interact via their (separable) inner product. Further, it is sufficient to assume that \mathcal{G} is fully separable, $\mathcal{G}(w) = \sum_j G(w_j)$, since each latent weight θ_j can be mapped to its binarized surrogate w_j independently (an underlying assumption in the majority of works but explicitly deviated from in [17]). Thus, the general form for \mathcal{E} assumed in the following is given by

$$\mathcal{E}(w; \theta) = \sum_j (G(w_j) - w_j \theta_j). \quad (13)$$

Therefore in this setting the solution $w^* = (w_1^*, \dots, w_d^*)^\top$ is given element-wise,

$$w_j^* = \arg \min_{w_j} G(w_j) - w_j \theta_j. \quad (14)$$

4.2. Relaxing by Optimal Value Reformulation

The optimal value reformulation (e.g. [31, 48]), which is a commonly used reformulation approach in bilevel optimization, allows us to rewrite the bilevel problem (9) as follows,

$$\min_{\theta, w} \ell(w) \quad \text{s.t.} \quad \mathcal{E}(w; \theta) \leq \min_{w'} \mathcal{E}(w'; \theta). \quad (15)$$

Observe that the w^* in the outer objective of (9) was replaced by a new unknown w , while the difficult equality constraint in (9) has been replaced by a somewhat easier inequality constraint. Due to the separable nature of \mathcal{E} in (13), it is advantageous to introduce an inequality constraint for each element w_j . Thus, we obtain

$$\min_{\theta, w} \ell(w) \quad \text{s.t.} \quad E(w_j; \theta_j) \leq \min_{w'_j} E(w'_j; \theta_j), \quad (16)$$

where E (independent of j) is given as

$$E(w_j; \theta_j) := G(w_j) - w_j \theta_j. \quad (17)$$

This first step enables us to straightforwardly relax (16) by fixing positive Lagrange multipliers for the inequality constraints:

$$\min_{\theta, w} \ell(w) + \sum_j \frac{1}{\beta_j} (E(w_j; \theta_j) - \min_{w'_j} E(w'_j; \theta_j)). \quad (18)$$

We parametrize the non-negative multipliers via β_j^{-1} for $\beta_j > 0$, which will be convenient in the following. Since we are interested in gradient-based methods, we replace the typically highly non-convex ‘‘loss’’ ℓ (which subsumes the target loss and the mapping induced by the network) by its linearization at w^* , $\ell(w^*) + (w - w^*)^\top \ell'(w^*)$. Recall that $w^* = \arg \min_w \mathcal{E}(w; \theta)$ is the effective weight used in the DNN and is ideally close to $\text{s\grave{g}n}(\theta)$. Overall, we arrive at the following relaxed objective to train a network with binary weights:

$$\begin{aligned} \mathcal{L}(\theta) &= \ell(w^*) - (w^*)^\top \ell'(w^*) \\ &+ \sum_j \min_{w_j} \left\{ w_j \ell'_j(w^*) + \frac{1}{\beta_j} E(w_j; \theta_j) \right\} \\ &- \sum_j \min_{w_j} \left\{ \frac{1}{\beta_j} E(w_j; \theta_j) \right\}, \end{aligned} \quad (19)$$

Using a linearized target loss above will be connected to a perturbed chain rule in Section 4.3. The inner minimization problems have the solutions

$$\begin{aligned} w_j^* &= \arg \min_{w_j} E(w_j; \theta_j) \quad \text{and} \\ \hat{w}_j &:= \arg \min_{w_j} \beta_j \ell'_j(w^*) w_j + E(w_j; \theta_j). \end{aligned} \quad (20)$$

$\hat{w} = (\hat{w}_1, \dots, \hat{w}_d)^\top$ is based on a perturbed objective that incorporates the local (first-order) behavior of the outer loss ℓ . Both w^* and \hat{w} implicitly depend on the current value of θ , and \hat{w} depends on a chosen ‘‘step size’’ vector $\vec{\beta} := (\beta_j)_{j=1}^d$ with each $\beta_j > 0$. If $\mathcal{E}(\cdot; \theta)$ is continuous at $w = w^*$, then $\lim_{\beta_j \rightarrow 0^+} \hat{w}_j = w_j^*$. Further, if \mathcal{E} is of the form given in (12), then \hat{w} is as easy to compute as w^* (proof in the supplementary material):

Proposition 1. *Let $\mathcal{E}(w; \theta) = G(w) - w^\top \theta$ and $w^* = \arg \min_w \mathcal{E}(w; \theta)$ be explicitly given as $w^* = \vec{s}(\theta)$. Then*

$$\hat{w} = \vec{s}(\theta - \vec{\beta} \odot \ell'(w^*)). \quad (21)$$

All of the interesting choices \mathcal{E} lead to efficient forward mappings s (like the choices $\mathcal{E}_{\text{tanh}}$ and $\mathcal{E}_{\text{hard-tanh}}$ given earlier that resulted in tanh and hard tanh functions).

4.3. Updating the latent weights θ

For a fixed choice of $\vec{\beta} = (\beta_1, \dots, \beta_d)^\top$ with $\beta_j > 0$, the relaxed objective $\mathcal{L}(\theta)$ in (19) is a nested minimization instance with a ‘‘min-min-max’’ structure. In some cases it is possible to obtain a pure ‘‘min-min-min’’ instance via

duality [49], but in practice this is not necessary. Let $\theta^{(t)}$ be the current solution at iteration t , then our employed local model to determine the new iterate $\theta^{(t+1)}$ is given by

$$Q(\theta; \theta^{(t)}) = \sum_j \frac{1}{\beta_j} (E(\hat{w}_j; \theta_j) - E(w_j^*; \theta_j)) + \frac{1}{2\eta} \|\theta - \theta^{(t)}\|^2, \quad (22)$$

where $w^* = \bar{s}(\theta^{(t)})$ and $\hat{w} = \bar{s}(\theta^{(t)} - \vec{\beta} \odot \ell'(w^*))$ are the effective weights and its perturbed instance, respectively, evaluated at $\theta^{(t)}$. The last term in Q regularizes deviations from $\theta^{(t)}$, and η plays the role of the learning rate. Minimizing $Q(\theta; \theta^{(t)})$ w.r.t. θ yields a gradient descent-like update,

$$\theta^{(t+1)} = \arg \min_{\theta} Q(\theta; \theta^{(t)}) = \theta^{(t)} - \eta(w^* - \hat{w}) \odot \vec{\beta} \quad (23)$$

for the assumed form of \mathcal{E} in (12). Each element of $(w^* - \hat{w}) \odot \vec{\beta}$, i.e. $(w_j^* - \hat{w}_j)/\beta_j$, corresponds to a finite difference approximation (using backward differences) of

$$-\frac{d}{d\beta_j} s(\theta_j^{(t)} - \beta_j \ell'_j(w^*)) \Big|_{\beta_j=0^+} \quad (24)$$

with spacing parameter $h_j = \beta_j \ell'_j(w^*)$. If s is at least one-sided differentiable, then it can be shown that these finite differences converge to a derivative given by the chain rule when $\beta_j \rightarrow 0^+$ [48],

$$\begin{aligned} \frac{1}{\beta_j} (w_j^* - \hat{w}_j) \xrightarrow{\beta_j \rightarrow 0^+} -\frac{d}{d\beta} s(\theta_j^{(t)} - \beta_j \ell'_j(w^*)) \Big|_{\beta_j=0^+} \\ = \ell'_j(s(\theta_j^{(t)})) s'(\theta_j^{(t)}) = \frac{d}{d\theta_j} \ell(s(\theta^{(t)})). \end{aligned} \quad (25)$$

For non-infinitesimal $\beta_j > 0$ the finite difference slope $(w_j^* - \hat{w}_j)/\beta_j$ corresponds to a perturbed chain rule,

$$\frac{1}{\beta_j} (w_j^* - \hat{w}_j) = \ell'_j(w^*) s'(\theta_j^{(t)} - \gamma_j \ell'_j(w^*)) \quad (26)$$

(recall that $w^* = s(\theta^{(t)})$), where the inner derivative is evaluated at a perturbed argument $\theta^{(t)} - \vec{\gamma} \odot \ell'(w^*)$ for a $\vec{\gamma} \in [0, \vec{\beta}]$. This is a consequence of the mean value theorem. Moreover, if each β_j is a critical point of the mapping

$$\beta \mapsto \frac{1}{\beta} (w_j^* - \hat{w}_j) = \frac{1}{\beta} (w_j^* - s(\theta_j^{(t)} - \beta \ell'_j(w^*))), \quad (27)$$

then by using the quotient rule it is easy to see that $\vec{\gamma} = \vec{\beta}$, and therefore

$$\frac{1}{\beta_j} (w_j^* - \hat{w}_j) = \ell'_j(w^*) s'(\theta_j^{(t)} - \beta_j \ell'_j(w^*)). \quad (28)$$

Additionally, the relation in (26) can be interpreted as a particular instance of mirror descent (recall Sec. 3.2) as shown in the supplementary material. Overall, the above means that we can relatively freely select where s' is actually evaluated. Since s is naturally a ‘‘squashing’’ function mapping

\mathbb{R} to the bounded interval $[-1, 1]$, gradient-based training using s' usually suffers from the vanishing gradient problem. Using the relaxed reformulation for bilevel programs allows us to select β_j to obtain a desired descent direction as it will be described in Section 4.5.

The resulting gradient-based training method is summarized in Alg. 1. The algorithm is stated as full batch method, but the extension to stochastic variants working with mini-batches drawn from p_{data} is straightforward. In the following section we discuss our choice of \mathcal{E} and how to select suitable spacing parameters $\vec{\beta}^{(t)} > 0$ in each iteration. Since $\vec{\beta}^{(t)}$ is chosen adaptively based on the values of $\theta^{(t)}$ and $\ell'(w^*)$ and used to perturb the chain rule, we call the resulting algorithm the *adaptive straight-through estimator* (AdaSTE) training method. In the supplementary material we highlight its relations with ProxQuant and mirror descent training (and also discuss convergence properties).

Algorithm 1 AdaSTE training method.

- 1: Initialize $\theta^{(0)}$, choose learning rates $\eta^{(t)}$, $t = 1, \dots$
 - 2: **for** $t = 1, \dots$ **do**
 - 3: $w^* \leftarrow \bar{s}(\theta^{(t)})$
 - 4: Run regular back-propagation to determine $\ell'(w^*)$
 - 5: Determine $\vec{\beta}^{(t)}$ using (38)
 - 6: $\hat{w} \leftarrow \bar{s}(\theta^{(t)} - \vec{\beta}^{(t)} \odot \ell'(w^*))$
 - 7: $\theta^{(t+1)} \leftarrow \theta^{(t)} - \eta^{(t)}(w^* - \hat{w}) \odot \vec{\beta}^{(t)}$
 - 8: **end for**
-

4.4. Our choice for the inner objective \mathcal{E}

In this section we will specify our choice for \mathcal{E} (and thus the mapping $\bar{s} : \theta \mapsto \arg \min_w \mathcal{E}(w; \theta)$). The straightforward options of $\mathcal{E}_{\text{tanh}}$ and $\mathcal{E}_{\text{hard-tanh}}$ (Section 4.1) suffer from the fact that the induced arg-min mappings coincide exactly with the sign function only when the hyper-parameter $\tau^{-1} \rightarrow \infty$. We are interested in an inner objective \mathcal{E} that yields perfect quantized mappings for finite-valued choices of hyper-parameters. Inspired by the double-well cost used in ProxQuant [6], we design \mathcal{E} as follows,

$$\mathcal{E}(w; \theta) = \frac{1+\mu}{2} \|w\|^2 - w^\top \theta - \mu(1+\alpha) \|w\|_1 + \iota_{[-1,1]^d}(w), \quad (29)$$

where $\mu > 0$ and $\alpha \in (0, 1)$ are free parameters. Note that \mathcal{E} is only piecewise convex in w for fixed θ , but it is fully separable in w_j with

$$E(w_j; \theta_j) = \frac{1+\mu}{2} w_j^2 - w_j \theta_j - \mu(1+\alpha) |w_j| + \iota_{[-1,1]}(w_j). \quad (30)$$

Via algebraic manipulations we find the following closed-form expression for \hat{w}_j (where we abbreviate ℓ' for $\ell'(w^*)$),




Figure 2. The graph of the mapping $w^* = s(\theta)$ given in (31) for $\alpha = 1/100$ and three different values of μ .

$$\begin{aligned} \hat{w}_j &= \arg \min_{w_j} \beta_j \ell'_j w_j + E(w_j; \theta_j) \\ &= \Pi_{[-1,1]} \left(\frac{\tilde{\theta}_j + \mu(1 + \alpha) \operatorname{sgn}(\tilde{\theta}_j)}{1 + \mu} \right), \end{aligned} \quad (31)$$

with $\tilde{\theta}_j := \theta_j - \beta_j \ell'_j$. In other words, the forward mapping $\vec{s}: \theta \mapsto w^* = s(\theta)$ for our choice of \mathcal{E} is given by

$$\vec{s}(\theta) = \Pi_{[-1,1]^d} \left(\frac{\theta + \mu(1 + \alpha) \operatorname{sgn}(\theta)}{1 + \mu} \right). \quad (32)$$

The piece-wise linear graph of this mapping is illustrated in Fig. 2 for $\alpha = 1/100$ and three different choices of μ . Let $\alpha \in (0, 1)$ be given, then $\vec{s}(\theta)$ attains only values in $\{-1, 1\}^d$ even for finite μ , since

$$\begin{aligned} \frac{|\theta_j| + \mu(1 + \alpha)}{1 + \mu} \geq 1 &\iff |\theta_j| + \mu(1 + \alpha) \geq 1 + \mu \\ \iff |\theta_j| + \alpha\mu &\geq 1, \end{aligned} \quad (33)$$

which implies that any θ_j is always mapped to +1 or -1 when $\mu \geq 1/\alpha$ (and the exact values of μ and α do not matter in this case). Consequently we have both the option to train with strictly binary weights from the beginning, or to train via a homotopy method by adjusting α or μ . Both choices lead to competitive results with the homotopy-based method having a small advantage in some cases as demonstrated in Section 5.

4.5. Adaptive choice for β

As indicated in Section 4.3, we can steer the modified chain rule by selecting $\beta_j > 0$ appropriately in order to determine a suitable descent direction. Note that each element θ_j in the vector of parameters θ has its own value for β_j . Below we describe how β_j is chosen when α and μ satisfy $\mu\alpha \geq 1$. In this setting we always have $w_j^* = \operatorname{sgn}(\theta_j) \in \{-1, 1\}$ and $\hat{w}_j = \operatorname{sgn}(\theta_j - \beta_j \ell'_j(w^*)) \in \{-1, 1\}$ (we ignore the theoretical possibility of $\theta_j = 0$ or $\theta_j - \beta_j \ell'_j(w^*) = 0$). Our aim is to select $\beta_j > 0$ such that the slope induced by backward differences, $\frac{1}{\beta_j}(w_j^* - \hat{w}_j)$,

is as close to $\ell'_j(w^*)$ as possible. In the following we abbreviate $\ell'(w^*)$ to ℓ' . Since sgn is an increasing step-function with derivative being zero almost everywhere, its finite difference approximation

$$\frac{1}{\beta_j}(w_j^* - \hat{w}_j) = \frac{1}{\beta_j}(\operatorname{sgn}(\theta_j) - \operatorname{sgn}(\theta_j - \beta_j \ell'_j)) \quad (34)$$

lies either in the interval $[0, s_{\max}]$ or in $[-s_{\max}, 0]$ for a suitable $s_{\max} \geq 0$ (which is dependent on θ_j and ℓ'_j). In particular, if $\theta_j \ell'_j \leq 0$, then $\operatorname{sgn}(\theta_j) = \operatorname{sgn}(\theta_j - \beta_j \ell'_j)$ for all $\beta_j \geq 0$ and $s_{\max} = 0$. On the other hand, if $\theta_j \ell'_j > 0$, then $\operatorname{sgn}(\theta_j - \beta_j \ell'_j) \neq \operatorname{sgn}(\theta_j)$ for $\beta_j > \theta_j / \ell'_j$ and therefore

$$\sup_{\beta_j > \theta_j / \ell'_j} \frac{1}{\beta_j} |w_j^* - \hat{w}_j| = \frac{2\ell'_j}{\theta_j}. \quad (35)$$

If θ_j is close to 0, then the r.h.s. may grow arbitrarily large (reflecting the non-existence of the derivative of sgn at 0). Assuming that $(w_j^* - \hat{w}_j)/\beta_j$ should maximally behave like a straight-through estimator (i.e. $|w_j^* - \hat{w}_j|/\beta_j \leq |\ell'_j|$, which also can be seen as a form of gradient clipping), we choose

$$\beta_j = \frac{1}{|\ell'_j|} \max\{2, |\theta_j|\} \quad (36)$$

in order to guarantee that

$$\frac{1}{\beta_j} |w_j^* - \hat{w}_j| \leq \frac{2}{\beta_j} \leq \frac{2|\ell'_j|}{2} = |\ell'_j|. \quad (37)$$

Overall, we obtain the following simple rule to assign each β_j for given θ and ℓ' :

$$\beta_j \leftarrow \begin{cases} \frac{1}{|\ell'_j|} \max\{2, |\theta_j|\} & \text{if } \theta_j \ell'_j > 0 \\ 1 & \text{otherwise.} \end{cases} \quad (38)$$

The choice of $\beta_j = 1$ in the alternative case is arbitrary, since $(w_j^* - \hat{w}_j)/\beta = 0$ for all values $\beta > 0$. Observe that the assignment of β_j in (38) selectively converts $(w_j^* - \hat{w}_j)/\beta_j$ into a scaled straight-through estimator whenever $\theta_j \ell'_j > 0$, otherwise the effective gradient used to update θ_j is zero (in agreement with the chain rule).

In the supplementary material we discuss the setting $\mu\alpha < 1$, which yields in certain cases different expressions for β_j . Nevertheless, we use (38) in all our experiments.

5. Experimental Results

In this section, we show several experimental results to validate the performance of our proposed method and compare it against existing algorithms that achieve state-of-the-art performance for our particular problem settings. As mentioned above, we only consider the training of networks with fully binarized weights and real-valued activations.

Implementation	CIFAR-10		CIFAR-100		TinyImageNet
	VGG-16	ResNet-18	VGG-16	ResNet-18	ResNet-18
Full-precision (†)	93.33	94.84	71.50	76.31	58.35
BinaryConnect (*)	89.75±0.26	91.92±0.23	54.61±2.37	68.67±0.7	-
BinaryConnect (†)	89.04	91.64	59.13	72.14	49.65
ProxQuant(†)	90.11	92.32	55.10	68.35	49.97
PMF(†)	91.40	93.24	64.71	71.56	51.52
MD-softmax (†)	90.47	91.28	56.25	68.49	46.52
MD-softmax-s (†)	91.30	93.28	63.97	72.18	51.81
MD-softmax-s (*)	91.39±0.30	93.10±0.17	64.42±0.37	71.87±0.25	-
MD-tanh (†)	91.64	92.27	61.31	72.13	54.62
MD-tanh-s (†)	91.53	93.18	61.69	72.18	52.32
MD-tanh-s (*)	91.40±0.30	93.23±0.15	62.93±0.60	71.96±0.18	-
BayesBiNN (*)	90.68±0.07	92.28±0.09	65.92±0.18	70.33±0.25	54.22
AdaSTE (w/o annealing) (*)	92.16±0.16	93.96±0.14	68.46±0.18	73.90±0.20	53.49
AdaSTE (with annealing) (*)	92.37±0.09	94.11±0.08	69.28±0.17	75.03±0.35	54.92

Table 1. Classification (test) accuracy for different methods. (*) indicates that experiments have been run 5 times using different random seeds (except for TinyImageNet). (†) indicates that results are obtained from the numbers reported by [3].

Following previous works [3, 6, 29], we use classification as the main task throughout our experiments. In particular, we evaluate the performance of the algorithms on the two network architectures: ResNet-18 and VGG16. The networks are trained and evaluated on the CIFAR10, CIFAR100 and TinyImageNet200 [25] datasets. We compare our algorithm against state-of-the-art approaches, including BinaryConnect (BC) [10], ProxQuant (PQ) [6], Proximal Mean-Field (PMF) [2], BayesBiNN [29], and several variants of Mirror Descent (MD) [3]. We employ the same standard data augmentations and normalization as employed by the methods we compare against (please refer to our supplemental material for more details about the experimental setup). Our method² is implemented in Pytorch and is developed based on the software framework released by BayesBiNN’s authors³ (more details regarding our implementation and additional Imagenette [20] results can be found in the supplemental material).

5.1. Classification Accuracy

In Table 1, we report the testing accuracy obtained by the considered methods. For PQ, PMF, the unstable versions of MD as well as for full-precision reference networks, we use the test accuracy reported in [3]. For BC, the stable variants of MD (i.e. MD-softmax-s and MD-tanh-s), we reproduce the results by running the source code released by the authors⁴ (using the default recommended hyper-parameters) for 5 different random initializations, and reporting the mean and standard deviation ob-

tained from these runs. The same strategy is also applied to BayesBiNN (hyper-parameters for BayesBiNN can be found in the supplemental material), except for the TinyImageNet dataset where we only report results for a single run (due to longer training time). We report the results for our method using two settings:

- Without annealing: we fix $\alpha = 0.01$ and $\mu = \frac{1}{\alpha}$.
- With annealing: we also use $\alpha = 0.01$ and set the initial value μ to $\mu^{(0)} = 1.0$, then increase μ after each epoch by a factor of γ , i.e. $\mu^{(t)} \leftarrow \gamma\mu^{(t-1)}$. γ is chosen such that μ reaches $1/\alpha$ after ≈ 200 epochs.

The impact of the choice of μ on the shape of $\vec{s}(\theta)$ is illustrated in Fig. 2. Table 1 demonstrates that our proposed algorithm achieves state-of-the-art results. Note that we achieve highly competitive results even without annealing μ (although annealing improves the test accuracy slightly but consistently). Hence, we conclude that AdaSTE without annealing (and therefore no additional hyper-parameters) can be used as a direct replacement for BinaryConnect. Note that we report all results after training for 500 epochs. In the supplementary material, we will show that both BayesBiNN and AdaSTE yield even higher accuracy if the models are trained for higher number of epochs.

5.2. Evolution of Loss and Accuracy

We further investigate the behavior of the algorithms during training. In particular, the evolution of training losses and testing accuracies are of interest, since these quantities are of practical interest.

In Fig. 3, we plot the testing accuracy obtained by our method in comparison with BC, MD (using the tanh map-

²<https://github.com/intellhave/AdaSTE>

³<https://github.com/team-approx-bayes/BayesBiNN>

⁴<https://github.com/kartikgupta-at-anu/md-bnn>




Figure 3. Testing accuracy achieved by the methods for the first 200 epochs with ResNet-18 (left) VGG16 (right) for CIFAR10 dataset (plots for CIFAR100 can be found in the supplemental material).




Figure 4. Training loss of the methods for the first 200 epochs with ResNet-18 (left) and VGG16 (right) on the CIFAR10 dataset (see supplementary material for plots of the CIFAR100 dataset).

ping), and BayesBiNN for the first 200 epochs. For our method, we show the performance for both settings with and without annealing (as described earlier). To obtain the plots for MD and BayesBiNN, we use the code provided by the authors with the default recommended hyper parameters. For BC, we use the implementation provided by the MD authors. As can be observed, AdaSTE quickly reaches very high test accuracy compared to other approaches. The MD-tanh approach (using the recommended annealing schedule from the authors [3]) only reaches satisfactory accuracy after approximately 100 epochs. We also try starting MD-tanh with a larger annealing parameter (i.e. the β hyper-parameter in [3]), but that yields very poor results (see supplemental material for more details). AdaSTE, on the other hand, is quite insensitive to the annealing details, and yields competitive results even without annealing.

Fig. 4 depicts the training loss of our methods compared to BayesBiNN. We choose to compare AdaSTE against our main competitor, BayesBiNN, as we have full control of the source code to assure that both methods are initialized with the same starting points. As can be seen, our method quickly reduces the training loss, while BayesBiNN takes longer for the training loss to converge. Note that BayesBiNN leverages the reparametrization trick and relies therefore on weights sampled from respective distributions at training time. In that sense AdaSTE is a purely deterministic algorithm, and the only source of stochasticity is the sampled mini-batches. This might be a factor explaining

AdaSTE’s faster reduction of the training loss.

6. Discussion and Conclusion

In this work we propose AdaSTE, an easy-to-implement replacement for the straight-through gradient estimator, and we demonstrate its benefits for training DNNs with strictly binary weights. One clear limitation in this work is, that we focus on the binary weight but real-valued activations scenario, which is a highly useful setting, but still prevents low-level implementations using only xor and bit count operations. Extending AdaSTE to binary activations seems straightforward, but will be more difficult to justify theoretically, and we expect training to be more challenging in practice. One obvious further shortcoming is our restriction to purely binary quantization levels, in particular to the set $\{+1, -1\}$. Generalizing the approach to arbitrary quantization levels can be done in several ways, e.g. by extending the W-shaped cost E in (30) to more minima or by moving to higher dimensions (e.g. by modeling parameters in the probability simplex).

Since weight quantization is one option to regulate the Lipschitz property of a DNNs’ forward mapping (and also its expressive power), the impact of weight quantization [12, 39] (and more generally DNN model compression [15, 46]) on adversarial robustness has been recently explored. Hence, combining our adaptive straight-through gradient estimator with adversarial training is one direction of future work.

References

- [1] Jan Achterhold, Jan Mathias Koehler, Anke Schmeink, and Tim Genewein. Variational network quantization. In *International Conference on Learning Representations*, 2018.
- [2] Thalaiyasingam Ajanthan, Puneet K. Dokania, Richard Hartley, and Philip H. S. Torr. Proximal mean-field for neural network quantization. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, October 2019.
- [3] Thalaiyasingam Ajanthan, Kartik Gupta, Philip Torr, Richard Hartley, and Puneet Dokania. Mirror descent view for neural network quantization. In Arindam Banerjee and Kenji Fukumizu, editors, *Proceedings of The 24th International Conference on Artificial Intelligence and Statistics*, volume 130 of *Proceedings of Machine Learning Research*, pages 2809–2817. PMLR, 13–15 Apr 2021.
- [4] Milad Alizadeh, Javier Fernández-Marqués, Nicholas D. Lane, and Yarin Gal. A systematic study of binary neural networks’ optimisation. In *International Conference on Learning Representations*, 2019.
- [5] Anish Athalye, Nicholas Carlini, and David Wagner. Obfuscated gradients give a false sense of security: Circumventing defenses to adversarial examples. In Jennifer Dy and Andreas Krause, editors, *Proceedings of the 35th International Conference on Machine Learning*, volume 80 of *Proceedings of Machine Learning Research*, pages 274–283. PMLR, 10–15 Jul 2018.
- [6] Yu Bai, Yu-Xiang Wang, and Edo Liberty. Proxquant: Quantized neural networks via proximal operators. In *International Conference on Learning Representations*, 2019.
- [7] Amir Beck and Marc Teboulle. Mirror descent and nonlinear projected subgradient methods for convex optimization. *Operations Research Letters*, 31(3):167–175, 2003.
- [8] Yoshua Bengio, Nicholas Léonard, and Aaron C. Courville. Estimating or propagating gradients through stochastic neurons for conditional computation. *CoRR*, abs/1308.3432, 2013.
- [9] Jiasi Chen and Xukan Ran. Deep learning with edge computing: A review. *Proceedings of the IEEE*, 107(8):1655–1674, 2019.
- [10] Matthieu Courbariaux, Yoshua Bengio, and Jean-Pierre David. BinaryConnect: Training Deep Neural Networks with binary weights during propagations. In C Cortes, N Lawrence, D Lee, M Sugiyama, and R Garnett, editors, *Advances in Neural Information Processing Systems*, volume 28. Curran Associates, Inc., 2015.
- [11] Lei Deng, Guoqi Li, Song Han, Luping Shi, and Yuan Xie. Model compression and hardware acceleration for neural networks: A comprehensive survey. *Proceedings of the IEEE*, 108(4):485–532, 2020.
- [12] Kirsty Duncan, Ekaterina Komendantskaya, Robert Stewart, and Michael Lones. Relative robustness of quantized neural networks against adversarial attacks. In *2020 International Joint Conference on Neural Networks (IJCNN)*, pages 1–8. IEEE, 2020.
- [13] Steven K. Esser, Jeffrey L. McKinstry, Deepika Bablani, Rathinakumar Appuswamy, and Dharmendra S. Modha. Learned step size quantization. In *International Conference on Learning Representations*, 2020.
- [14] Tal Grossman. The CHIR Algorithm for Feed Forward Networks with Binary Weights. In D Touretzky, editor, *Advances in Neural Information Processing Systems*, volume 2. Morgan-Kaufmann, 1990.
- [15] Shupeng Gui, Haotao N Wang, Haichuan Yang, Chen Yu, Zhangyang Wang, and Ji Liu. Model compression with adversarial robustness: A unified optimization framework. *Advances in Neural Information Processing Systems*, 32:1285–1296, 2019.
- [16] Yunhui Guo. A survey on methods and theories of quantized neural networks. *CoRR*, abs/1808.04752, 2018.
- [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*, pages 4017–4026. PMLR, 2020.
- [18] 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 H Wallach, H Larochelle, A Beygelzimer, F d’ Alché-Buc, E Fox, and R Garnett, editors, *Advances in Neural Information Processing Systems*, volume 32. Curran Associates, Inc., 2019.
- [19] G Hinton. Neural networks for machine learning, coursera video lectures, 2012.
- [20] Jeremy Howard and Sylvain Gugger. Fastai: a layered api for deep learning. *Information*, 11(2):108, 2020.
- [21] Itay Hubara, Matthieu Courbariaux, Daniel Soudry, Ran El-Yaniv, and Yoshua Bengio. Binarized Neural Networks. In D Lee, M Sugiyama, U Luxburg, I Guyon, and R Garnett, editors, *Advances in Neural Information Processing Systems*, volume 29. Curran Associates, Inc., 2016.
- [22] Sergey Ioffe and Christian Szegedy. Batch normalization: Accelerating deep network training by reducing internal covariate shift. In Francis Bach and David Blei, editors, *Proceedings of the 32nd International Conference on Machine Learning*, volume 37 of *Proceedings of Machine Learning Research*, pages 448–456, Lille, France, 07–09 Jul 2015. PMLR.
- [23] Eric Jang, Shixiang Gu, and Ben Poole. Categorical reparameterization with gumbel-softmax. *arXiv preprint arXiv:1611.01144*, 2016.
- [24] Mohammad Emtiyaz Khan and Håvard Rue. The bayesian learning rule, 2021.
- [25] Ya Le and Xuan Yang. Tiny imagenet visual recognition challenge. *CS 231N*, 7(7):3, 2015.
- [26] Mingbao Lin, Rongrong Ji, Zihan Xu, Baochang Zhang, Yan Wang, Yongjian Wu, Feiyue Huang, and Chia-Wen Lin. Rotated Binary Neural Network. In H Larochelle, M Ranzato, R Hadsell, M F Balcan, and H Lin, editors, *Advances in Neural Information Processing Systems*, volume 33, pages 7474–7485. Curran Associates, Inc., 2020.
- [27] Christos Louizos, Matthias Reisser, Tijmen Blankevoort, Efstratios Gavves, and Max Welling. Relaxed quantization for discretized neural networks. In *International Conference on Learning Representations*, 2019.

- [28] C Maddison, A Mnih, and Y Teh. The concrete distribution: A continuous relaxation of discrete random variables. In *Proceedings of the international conference on learning Representations*. International Conference on Learning Representations, 2017.
- [29] Xiangming Meng, Roman Bachmann, and Mohammad Emteyaz Khan. Training binary neural networks using the Bayesian learning rule. In Hal Daumé III and Aarti Singh, editors, *Proceedings of the 37th International Conference on Machine Learning*, volume 119 of *Proceedings of Machine Learning Research*, pages 6852–6861. PMLR, 13–18 Jul 2020.
- [30] Arkadij Semenov and David Borsovich. *Problem complexity and method efficiency in optimization*. Wiley-Interscience, 1983.
- [31] Jirí V Outrata. A note on the usage of nondifferentiable exact penalties in some special optimization problems. *Kybernetika*, 24(4):251–258, 1988.
- [32] Haotong Qin, Ruihao Gong, Xianglong Liu, Xiao Bai, Jingkuan Song, and Nicu Sebe. Binary neural networks: A survey. *Pattern Recognition*, 105:107281, Sep 2020.
- [33] Haotong Qin, Ruihao Gong, Xianglong Liu, Mingzhu Shen, Ziran Wei, Fengwei Yu, and Jingkuan Song. Forward and backward information retention for accurate binary neural networks. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2020.
- [34] Mohammad Rastegari, Vicente Ordonez, Joseph Redmon, and Ali Farhadi. XNOR-Net: ImageNet Classification Using Binary Convolutional Neural Networks. In Bastian Leibe, Jiri Matas, Nicu Sebe, and Max Welling, editors, *Computer Vision – ECCV 2016*, pages 525–542, Cham, 2016. Springer International Publishing.
- [35] F. Rosenblatt. *The Perceptron, a Perceiving and Recognizing Automaton Project Para*. Report: Cornell Aeronautical Laboratory. Cornell Aeronautical Laboratory, 1957.
- [36] F. Rosenblatt. *Principles of Neurodynamics: Perceptrons and the Theory of Brain Mechanisms*. Cornell Aeronautical Laboratory. Report no. VG-1196-G-8. Spartan Books, 1962.
- [37] D Saad and E Marom. Training Feed Forward Nets with Binary Weights Via a Modified CHIR Algorithm. *Complex Systems*, 4:573–586, 1990.
- [38] Benjamin Scellier and Yoshua Bengio. Equilibrium Propagation: Bridging the Gap between Energy-Based Models and Backpropagation. *Frontiers in Computational Neuroscience*, 11, 2017.
- [39] Chang Song, Elias Fallon, and Hai Li. Improving adversarial robustness in weight-quantized neural networks. *arXiv preprint arXiv:2012.14965*, 2020.
- [40] Daniel Soudry, Itay Hubara, and Ron Meir. Expectation backpropagation: Parameter-free training of multilayer neural networks with continuous or discrete weights. *Advances in Neural Information Processing Systems*, 2(January):963–971, 2014.
- [41] Evgeny Strelakovsky and Daniel Cremers. Real-time minimization of the piecewise smooth Mumford-Shah functional. In *European conference on computer vision*, pages 127–141. Springer, 2014.
- [42] Stefan Uhlich, Lukas Mauch, Fabien Cardinaux, Kazuki Yoshiyama, Javier Alonso Garcia, Stephen Tiedemann, Thomas Kemp, and Akira Nakamura. Mixed precision dnns: All you need is a good parametrization. In *International Conference on Learning Representations*, 2020.
- [43] Santosh S. Venkatesh. Directed drift: A new linear threshold algorithm for learning binary weights on-line. *Journal of Computer and System Sciences*, 46(2):198–217, 1993.
- [44] Lin Xiao. Dual averaging methods for regularized stochastic learning and online optimization. *Journal of Machine Learning Research*, 11(88):2543–2596, 2010.
- [45] Xiaohui Xie and H. Sebastian Seung. Equivalence of back-propagation and contrastive Hebbian learning in a layered network. *Neural Computation*, 15(2):441–454, 2003.
- [46] Shaokai Ye, Kaidi Xu, Sijia Liu, Hao Cheng, Jan-Henrik Lambrechts, Huan Zhang, Aojun Zhou, Kaisheng Ma, Yanzhi Wang, and Xue Lin. Adversarial robustness vs. model compression, or both? In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 111–120, 2019.
- [47] 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.
- [48] Christopher Zach. Bilevel programming and deep learning: A unifying view on inference learning methods. *CoRR*, abs/2105.07231, 2021.
- [49] Christopher Zach and Virginia Estellers. Contrastive learning for lifted networks. In *British Machine Vision Conference*, 2019.