
ASkewSGD : An Annealed interval-constrained Optimisation method to train Quantized Neural Networks

Louis Leconte
LISITE, Sorbonne University
Mathematical and Algorithmic
Sciences Lab (Huawei)

Sholom Schechtman
CITI, Télécom SudParis

Eric Moulines
CMAP, Ecole Polytechnique

Abstract

In this paper, we develop a new algorithm, Annealed Skewed SGD - ASkewSGD - for training deep neural networks (DNNs) with quantized weights. First, we formulate the training of quantized neural networks (QNNs) as a smoothed sequence of interval-constrained optimization problems. Then, we propose a new first-order stochastic method, ASkewSGD, to solve each constrained optimization subproblem. Unlike algorithms with active sets and feasible directions, ASkewSGD avoids projections or optimization under the entire feasible set and allows iterates that are infeasible. The numerical complexity of ASkewSGD is comparable to existing approaches for training QNNs, such as the straight-through gradient estimator used in BinaryConnect, or other state of the art methods (ProxQuant, LUQ). We establish convergence guarantees for ASkewSGD (under general assumptions for the objective function). Experimental results show that the ASkewSGD algorithm performs better than or on par with state of the art methods in classical benchmarks.

1 Introduction

The use of deep neural networks (DNNs) on computing hardware such as mobile and IoT devices with limited computational and memory resources is becoming increasingly important. This has led to a growing area of research focused on reducing the model size and inference time of DNNs; in this area, the overall goal is to keep the loss of accuracy below an acceptable level compared to floating-

point implementations. These methods include, for example, model pruning, neural architecture search, novel efficient architecture design, and low-rank decomposition. In this work, we focus on network quantization, where weights and activations are quantized to lower bit widths, allowing for efficient fixed-point inference and reduced memory bandwidth usage; see, for example, Courbariaux et al. (2015); Jacob et al. (2018); Darabi et al. (2018); Choukroun et al. (2019); Deng et al. (2020); Qin et al. (2020); Bhalgat et al. (2020); Chmiel et al. (2021) and references therein. Quantized neural networks (QNNs) have attracted many research efforts. Nevertheless, the challenge of closing the accuracy gap between full-precision and quantized networks remains open, especially for extremely low-precision arithmetics (e.g. binary). The task of learning a quantized neural network (QNN) can be formulated as minimising the training loss with quantization constraints on the weights, i.e.,

$$\min_{w \in \mathcal{Q}} \ell(w), \ell(w) = \mathbb{E}_{(x,y) \sim p_{\text{data}}} [\ell(f(x, w), y)], \quad (1)$$

where $\mathcal{Q} \subset \mathbb{R}^d$ is the set of quantization levels, d is the number of parameters (network weights and biases), ℓ is the training loss (e.g. the cross-entropy or square loss), $f(x, w)$ is the DNN prediction function, p_{data} is the training distribution. The quantization constraints in the above program make it an extremely difficult task: the underlying optimization problem is non-convex, non-differentiable, and combinatorial in nature. Optimization of smooth functions of integer valued variables (and even quadratic ones like the max-cut problem in graph theory) is known to be NP-hard (Garey and Johnson, 1980). The challenge is to find algorithms that can produce a sensible approximate solution with a manageable computational effort. Inspired by mixed-integer nonlinear programming (MINLP) problems, several approaches using geometric, analytic, and algebraic techniques have been proposed to transform the discrete problem into a continuous problem. Examples include the use of global or concave optimization formulations, semidefinite programming, and spectral theory (see e.g. Mitchell et al. (1998); Bussieck et al. (2003); Horst and Tuy (2013); Beck and Teboulle (2000); Murray and Ng (2010)). However, these types of approaches are doomed to fail in the NN

context because the number of parameters is several orders of magnitude larger than for classical MINLP problems.

For large training data sets and number of variables d , stochastic gradient-based (first-order) methods for finding minimizers of (1) are often the only manageable option. Several methods have been proposed which transform the loss function (1) into a differentiable surrogate (with possibly an additional penalty term) to "favor" quantized solutions. The general approach is to introduce real-valued "latent" weights $w \in \mathbb{R}^d$ from which the quantized weights are generated; in the binary case, it is classical to use the $\text{sign}(\cdot)$ function or a differentiable surrogate thereof. The simplest method, called BinaryConnect (BC) (Courbariaux et al., 2015), is based on straight-through estimators (STE) that ignore the sign conversion in computing the gradient with respect to the latent weights w . BC reaches state-of-the-art performance on elementary classification tasks and is still a competitive baseline method for more sophisticated problems. Extensions of STE has also been used for more general QNN by Chmiel et al. (2021); Sun et al. (2020); Choi et al. (2018); Wang et al. (2019).

However, despite its success in NN inference, the STE method does not rely on solid theoretical groundings and may be shown to fail on simple low-dimensional benchmarks - even with convex objective functions; see Bai et al. (2018) and Section 5.3. We discuss this method and its recent improvements in the paragraph on related works (see below).

Contributions

- We replace the discrete optimization (1) by an annealed sequence of differentiable inequality constraints that converges to (1) when the annealing parameter goes to 0.
- We use a novel first-order algorithm proposed in Muehlebach and Jordan (2021) to solve the relaxed subproblems in the annealed sequence, leading to ASkewSGD. Unlike classical constrained optimization algorithms, including the projection method or sequential quadratic programming (Gill and Wong, 2012), this approach relies exclusively on local approximations of the feasible set. This local approximation includes only the active constraints, and is guaranteed to be a convex polyhedron even if the underlying constraint set is non convex. This makes the resulting algorithm easy to implement and also ensures that the descent is not stopped as soon as a new constraint is violated.
- We show how ASkewSGD can be applied to train QNN. The complexity of the resulting algorithm is similar to that of BC or LUQ (Courbariaux et al., 2015; Chmiel et al., 2021) and ProxQuant (Bai et al., 2018). Our algorithm uses high precision latent weights and uses classical backpropagation to evaluate the gradients.
- We provide convergence guarantees for ASkewSGD. We

stress that, as opposed to Muehlebach and Jordan (2021), no convexity assumption on the objective function or the feasible set is made.

- We evaluate the performance of ASkewSGD on classical computer vision datasets using ConvNets and ResNets. Our experiments show that QNNs trained with ASkewSGD achieve accuracy very close to that of their floating-point counterparts, and outperform or are on par with comparable baselines.

2 Related works

We focus on BNN and QNN that replace floating-point multiplication and addition operations with efficient fixed-point arithmetic. We do not consider algorithms that use low-bit computations at the learning stage; see Sakr and Shanbhag (2018); Chen et al. (2020). Given the abundance of works, it is impossible to give complete references. We focus mostly on methods used in our benchmarks.

Binary NN: The first attempt to train BNN is BinaryConnect (BC) (Courbariaux et al., 2015; Hubara et al., 2016) which is the first algorithm to implement Quantization Aware Training (QAT); see (Gholami et al., 2021; Zhao et al., 2020; Guo, 2018; Nagel et al., 2021) and the references therein. BC uses full precision latent weights. On the forward path, the latent weights are binarized. On the backward path, classical backpropagation is applied to update the latent weights, using a differentiable proxy of the binarization function in the gradient calculation. The most common implementation uses the identity proxy, resulting in the straight-through estimator (STE). Although the neural network parameters are highly compressed (and quantization errors can be large), the BC-STE estimator and its numerous recent improvements perform satisfactorily in many benchmarks and have become a de facto standard; see Hu et al. (2018); Faraone et al. (2018); Le et al. (2021); Anderson and Berg (2018).

ProxQuant (PQ) (Bai et al., 2018), Proximal Mean-Field (PMF) (Ajanthan et al., 2019), Mirror Descent (MD) (Ajanthan et al., 2021), and Rotated Binary Neural Networks (RBNN) (Lin et al., 2020) formulate the task of training BNNs as a constrained optimization problem and discuss different methods to generate binary weights from real-valued latent weights. All of these methods have in common that they use gradual annealing of the conversion mapping, in the sense that, unlike BC and its variants, the latent weights are not projected onto a finite set of quantization values in the forward path. Instead, a force is applied to gradually push the latent weights to the quantization constraints, in a manner reminiscent of homotopy methods for solving nonlinear systems or penalty barrier in nonlinear optimization.

BNN as Variational Inference (VI): Training binary neural networks can also be approached with VI; see among

others Raiko et al. (2015); Peters and Welling (2018); Roth et al. (2019). Instead of optimizing binary weights, the parameters of Bernoulli distributions are learned using the VI Bayesian learning rule; see e.g. Khan and Rue (2021). Even if unbiased estimators of the ELBO are available, classical methods like MuProp (Gu et al., 2016) or REINFORCE with variance-reduction baselines (Mnih and Gregor, 2014) have a prohibitively high variance. The use of Gumbel-Softmax (GS) trick (Jang et al., 2016; Maddison et al., 2017) has been advocated in Meng et al. (2020), but as noted in (Shekhovtsov, 2021, Section 4) there is an issue in the implementation which paradoxically enables the training. The connections between STE algorithms and their many variants - including MD - and VI methods are further discussed in Shekhovtsov and Yanush (2021).

Quantized NN: The STE estimator is easily adapted to QNN by adding a projection step onto the set of quantization levels in the forward pass (Zhou et al., 2016); see (Choi et al., 2018; Sun et al., 2020; Chmiel et al., 2021) and the references therein. To mitigate performance loss reported in early work from Zhou et al. (2016), a number of attempts has been proposed. One possible way is to increase the NN size (Zagoruyko and Komodakis, 2016), or the number of channel for convolution layers (Mishra et al., 2017; McDonnell, 2018). Knowledge distillation has also been considered with some success (Mishra and Marr, 2017). A teacher network (typically very large (Liu et al., 2020) and trained in full-precision) is employed to help the QNN training (the student network).

In QNN, the choice of the quantizer and the normalization of the weights (at each layer) play a key role. Many works have been devoted to the design of non-uniform or statistical (distribution dependent) quantizers; see (Banner et al., 2018; Hou and Kwok, 2018; Bhalgat et al., 2020; Liang et al., 2021; Fournarakis and Nagel, 2021; Zhou et al., 2017, 2018) and the references therein. Statistical quantizers are often more efficient, but they are more complex to implement and often require fine tuning (Zhang et al., 2021).

A number of works have considered formulating the quantization problem as an optimization problem (Li et al., 2017, 2016; Zhu et al., 2016; Carreira-Perpinán and Idelbayev, 2017; Leng et al., 2018; Polino et al., 2018), but the proposed methods rely on assumptions which may not hold for deep neural networks (Guo, 2018). In Moons et al. (2017); Yang et al. (2017); Esser et al. (2015), the QNN training is tackled as an energy efficiency problem, whereas Gong et al. (2019) propose a Differentiable Soft Quantization (DSQ) to efficiently train QNN.

Activation function Quantization: We have so far described the quantization of the network weights. But an efficient implementation also requires the quantization of the activation functions. For BNN, (Kim and Smaragdus, 2016; Hubara et al., 2016; Rastegari et al., 2016) proposed

to use $\text{sign}(\cdot)$ function, but this approach significantly affects the performance. More complex quantization schemes have been considered in Choi et al. (2018) alleviating performance degradation. Hybrid formats FP8 (Wang et al., 2018) or INT8 (Wiedemann et al., 2020; Banner et al., 2018) were successfully employed to achieve a low precision training. Recent works have proposed to jointly optimize the quantization parameters (of weights and activations) and the weights parameters. This task can be done by modifying the learning loss or by minimizing the quantization error (Zhu et al., 2016; Zhang et al., 2018; Li et al., 2019).

3 Algorithm derivation

In this section we first introduce the Muehlebach and Jordan (2021) (MJ) algorithm for smooth constrained optimization, initially proposed in a convex setting. We describe the algorithm in full generality and then show how to adapt the MJ algorithm to the QNN setting.

The MJ algorithm Consider the following optimization problem:

$$\min_{w \in C} \ell(w), \quad C = \{w \in \mathbb{R}^d : g(w) \geq 0\}, \quad (2)$$

where $\ell : \mathbb{R}^d \rightarrow \mathbb{R}$ denotes the objective function, $g : \mathbb{R}^d \rightarrow \mathbb{R}^{n_g}$ define the inequality constraints. We assume that the feasible set C is non-empty and compact and that the functions ℓ and g are continuously differentiable. We stress that neither ℓ nor C are assumed to be convex. Standard solutions to find a local minimizer of (2) use either a projected gradient descent algorithm or “non-linear” projection like mirror descent. However, C might have a complicated form, in which case computing the projection on C might require to solve a non-trivial optimization algorithm in itself (and may fail to be properly defined). The basic idea behind Muehlebach and Jordan (2021)’s proposal is to “skew” the search direction in order to force the algorithm to find a minimizer of (2) without constraining the sequence $(w_k)_{k \in \mathbb{N}}$ to the feasible set. For any $w \in \mathbb{R}^d$, define by $I(w)$ the set of active constraints

$$I(w) = \{i \in \{1, \dots, n_g\}, g_i(w) \leq 0\}.$$

Under mild assumptions (basically, Muehlebach and Jordan (2021) assume that Mangasarian Fromowitz constraint qualification conditions hold everywhere and not on the feasible set only) the tangent and normal cones of C at $w \in C$ are given by:

$$\begin{aligned} T_C(w) &= \{v \in \mathbb{R}^d, \nabla g_i(w)^\top v \geq 0, \text{ for all } i \in I(w)\}, \\ N_C(w) &= \left\{ - \sum_{i \in I(w)} \lambda_i \nabla g_i(w), \lambda_i \in \mathbb{R}_+ \right\}. \end{aligned}$$

Moreover, the Karush-Kuhn-Tucker (KKT) conditions hold (Borwein and Lewis, 2006, Theorem 7.2.9): if w^* is a local




Figure 1: The vector field of velocities for $\epsilon = 0.3$, and $\alpha = 0.1$ (left panel) or $\alpha = 1.0$ (right panel). Here, $\ell(w^1, w^2) = ((w^1 - 0.5)^2 + (w^2 - 0.5)^2)/3$ and $g(w^1, w^2)^\top = (\epsilon - ((w^1)^2 - 1)^2, \epsilon - ((w^2)^2 - 1)^2)$. The border of the set of constraints is shown in blue, and the minimizer of the constrained and unconstrained optimization problems are shown with blue and green dots, respectively.

minimizer of (2), then $w^* \in \mathcal{Z}$, where

$$\mathcal{Z} := \{w \in C : 0 \in -\nabla\ell(w) - N_C(w)\}.$$

The MJ algorithm (Muehlebach and Jordan, 2021) generates iterates in \mathbb{R}^d as follows:

$$\begin{cases} w_{k+1} = w_k + \gamma_k v_k \\ v_k = \arg \min_{v \in V_\alpha(w_k)} (1/2) \|v + \nabla\ell(w_k)\|^2, \end{cases} \quad (3)$$

where (γ_k) is a non-increasing sequence of positive step sizes, $\alpha > 0$ is an hyper-parameter, and the sets $V_\alpha(w)$ are defined as:

$$V_\alpha(w) = \{v \in \mathbb{R}^d : \nabla g_i(w)^\top v \geq -\alpha g_i(w) \text{ for all } i \in I(w)\}.$$

If $w \in C$ and $i \in I(w)$, then $g_i(w) = 0$ and thus $V_\alpha(w)$ reduces to $T_C(w)$. The set $V_\alpha(w)$ can be considered as an extension of the tangent cone "outside" of the feasible set. Note also that $V_\alpha(w)$, for all $w \in \mathbb{R}^d$, is a convex polyhedron whose construction includes only the active constraints.

By construction, whenever $g_i(w_k) \leq 0$, $\nabla g_i(w_k)^\top v_k \geq -\alpha g_i(w_k)$. Thus, in Eq. (3), the velocity v_k is chosen to match the unconstrained gradient flow $-\nabla\ell(w_k)$ as closely as possible, subject to the velocity constraint $v_k \in V_\alpha(w_k)$ (this is illustrated on a simple example in Figure 1, for different values of $\alpha > 0$). A striking difference from the classical projected gradient algorithm is that the MJ approach is based on a local approximation of the feasible set. This local approximation includes only the active constraints and is guaranteed to be a convex polyhedron even if the underlying feasible set is not convex. In "classical" constrained optimization algorithms, constraints are typically handled by direct reference to positions, meaning that the iterates w_k , for all $k \geq 0$, must lie in the constraint set C .

ASkewSGD description In Muehlebach and Jordan (2021) the convergence of the MJ algorithm was proven

under the condition that the function ℓ and the set C are convex. We now adapt this algorithm to the QNN problem, removing the requirement that ℓ, C are convex and, furthermore, replacing $\nabla\ell$ by a mini-batch stochastic gradient.

We consider ℓ , the training loss, written as $\ell(w) = 1/N \sum_{j=1}^N \ell_j(w)$, where N is the size of the training set, and ℓ_j is the loss associated with the j -th observation.

We relax the quantization constraints $w^i \in \mathcal{Q}$, $i \in \{1, \dots, d\}$, to a sequence of "smoothed" interval constraints. The set of quantization values \mathcal{Q} is defined coordinate wise: $\{c_1^i, \dots, c_{K^i}^i\}$. We assume for full generality a different scalar quantizer for each coefficient; we typically use different scalar quantizers for each layer of the NN (but the same quantizer for the coefficients in the same layer). For $\omega \in \mathbb{R}$ such that $\omega \in [c_1^i, c_{K^i}^i]$, we define

$$\phi^i(\omega) = (\omega - c_{Q^i(\omega)}^i)^2 (\omega - c_{Q^i(\omega)+1}^i)^2$$

where $Q^i(\omega)$ is the unique index satisfying $c_{Q^i(\omega)}^i \leq \omega < c_{Q^i(\omega)+1}^i$. If $\omega < c_1^i$ we define $\phi^i(\omega) = (\omega - c_1^i)^2$, and if $\omega > c_{K^i}^i$ we define $\phi^i(\omega) = (\omega - c_{K^i}^i)^2$. For any $\epsilon \in [0, 1]$ and $\omega \in \mathbb{R}$, define $\psi_\epsilon^i(\omega) := \epsilon - \phi^i(\omega)$ and consider the feasible set

$$C_\epsilon = \{w \in \mathbb{R}^d : \forall i \in \{1, \dots, d\}, g_{\epsilon, i}(w) := \psi_\epsilon^i(w^i) \geq 0\}.$$

For each $\epsilon \in (0, 1)$, we consider the constrained optimization problem $\mathcal{P}_\epsilon : \min_{w \in C_\epsilon} \ell(w)$. It is easily seen that $\bigcap_{\epsilon > 0} C_\epsilon = \mathcal{Q}$, recovering the constraints of the QNN problem. We therefore define a decreasing sequence $(\epsilon_n)_{n \geq 0}$ of numbers in $[0, 1]$ such that $\lim_{n \rightarrow \infty} \epsilon_n = 0$ and solve (approximately) the sequence of problems $(\mathcal{P}_{\epsilon_n})_{n \in \mathbb{N}}$.

Here we must notice that the set $V_\alpha(w)$ is empty if and only if there is $1 \leq i \leq d$ such that $w^i = (c_{Q^i(\omega)}^i + c_{Q^i(\omega)+1}^i)/2$. For such a point, there is no "best" direction, so we chose it arbitrarily by specifying that the i -th coordinate must go to the right (see the following clipping convention). A symmetric choice prescribing a left direction is also possible. Moreover, since the set of such w is of Lebesgue-measure zero, we can hope that we will never stumble upon such a point (this is further guaranteed by the fact that the iterates converge to C_ϵ , implying that such points asymptotically never occur).

We denote by $\mathcal{Z}_\epsilon := \{w \in C_\epsilon : 0 \in -\nabla\ell(w) - N_{C_\epsilon}(w)\}$, the set of KKT points of \mathcal{P}_ϵ . Notice that any element of $N_{C_\epsilon}(w)$ can be written as $(-\lambda^1 \psi'_\epsilon(w^1), \dots, -\lambda^d \psi'_\epsilon(w^d))$, with $\lambda^i \geq 0$ and $\lambda^i \neq 0$ only if $\psi_\epsilon(w^i) = 0$. Therefore, $w \in \mathcal{Z}_\epsilon$ if and only if for every $i \in \{1, \dots, d\}$,

$$\begin{aligned} \nabla_i \ell(w) &= 0 && \text{if } \psi_\epsilon(w^i) > 0 \\ \text{and } \text{sign}(\nabla_i \ell(w)) &= \text{sign}(\psi'_\epsilon(w^i)) && \text{if } \psi_\epsilon(w^i) = 0, \end{aligned}$$

where for $i \in \{1, \dots, d\}$, $\nabla_i \ell(w)$ is the partial derivative of $\ell(w)$ w.r.t. w^i . In this setting, the set of active constraints

and V_α can be written down as:

$$\begin{aligned} \mathbf{I}_\epsilon(w) &= \{i \in \{1, \dots, d\} : \psi_\epsilon(w^i) \leq 0\}, \\ V_{\epsilon, \alpha}(w) &= \{v \in \mathbb{R}^d : v^i \psi'_\epsilon(w^i) \geq -\alpha \psi_\epsilon(w^i) \text{ for } i \in \mathbf{I}_\epsilon(w)\}. \end{aligned}$$

Let w be such that $w^i \neq (c_{Q^i(\omega)} + c_{Q^i(\omega)+1})/2$ for all $i \in \{1, \dots, d\}$. For $u \in \mathbb{R}^d$, denote by $s_{\epsilon, \alpha}(u, w) = \arg \min_{v \in V_{\epsilon, \alpha}(w)} 1/2 \|v + u\|^2$. This problem admits an explicit solution: $[s_{\epsilon, \alpha}(u, w)]^i = -u^i$ if $\psi_\epsilon(w^i) > 0$ or $-\psi'_\epsilon(w^i)u^i \geq -\alpha \psi_\epsilon(w^i) \geq 0$ and $[s_{\epsilon, \alpha}(u, w)]^i = -\alpha \psi_\epsilon(w^i) / \psi'_\epsilon(w^i)$, otherwise. Note that when $w^i \rightarrow (c_{Q^i(\omega)} + c_{Q^i(\omega)+1})/2$, the quantity $\psi'_\epsilon(w^i)$ converges to zero, and thus $[s_{\epsilon, \alpha}(u, w)]^i$ might diverge to infinity. To alleviate this problem, we furthermore clip the update. For $(a, b) \in \mathbb{R} \times \mathbb{R}_+$, define $\text{clip}(a, b)$ equal to a if $|a| \leq b$ and to $b \text{sign}(a)$ otherwise. Choose $M_{\epsilon, c} > 0$ and let $s_{\epsilon, \alpha}^c$ be defined for $i \in \{1, \dots, d\}$, $w^i \neq (c_{Q^i(\omega)} + c_{Q^i(\omega)+1})/2$, by: $[s_{\epsilon, \alpha}^c(u, w)]^i =$

$$\begin{cases} -u^i & \text{if } \psi_\epsilon(w^i) > 0 \text{ or } -\psi'_\epsilon(w^i)u^i \geq -\alpha \psi_\epsilon(w^i) \geq 0; \\ \text{clip}(-\alpha \psi_\epsilon(w^i) / \psi'_\epsilon(w^i), M_{\epsilon, c}) & \text{otherwise.} \end{cases}$$

We set by convention $[s_{\epsilon, \alpha}^c(u, w)]^i = M_{\epsilon, c}$ if $w^i = (c_{Q^i(\omega)} + c_{Q^i(\omega)+1})/2$. For given α, ϵ , ASkewSGD is summarized in Algorithm 1. Under mild assumptions, we estab-

Algorithm 1 ASkewSGD algorithm

Data: sequence of step sizes (γ_k) ; size of the mini-batch $N_b \leq N$; $w_0 \in \mathbb{R}^d$

- 1 **for** $k=1, \dots, T$ **do**
 - 2 Sample a minibatch of N_b observations $\{j_1, \dots, j_{N_b}\}$ in $\{1, \dots, N\}$;
 Compute the Stochastic Gradient $\widehat{\nabla} \ell(w_k) = 1/N_b \sum_{i=1}^{N_b} \nabla \ell_{j_i}(w_k)$;
 Compute the update direction $v_k = s_{\epsilon, \alpha}^c(\widehat{\nabla} \ell(w_k), w_k)$;
 Update the parameter $w_{k+1} = w_k + \gamma_k v_k$.
 - 3 **end**
-

lish the convergence of ASkewSGD. Consider the following assumptions.

A1. For $j \in \{1, \dots, N\}$, the function ℓ_j is d -times continuously differentiable and has M_{ℓ_j} -Lipschitz continuous gradients.

A2. The stepsizes $(\gamma_k)_{k \geq 0}$ are positive, $\sum_{j=0}^{\infty} \gamma_k = \infty$ and $\sum_{j=0}^{\infty} \gamma_k^2 < \infty$.

Notice that **A2** holds for (γ_k) of the form $(1/k^\delta)$, with $\delta \in (1/2, 1]$. **A1** will ensure the stability of ASkewSGD (i.e. the iterates are bounded with probability one). Moreover, **A1** implies that $\ell(\mathcal{Z}_\epsilon)$ is of empty interior, as a consequence of the Sard's theorem (see Lemma 5).

Theorem 1. Assume **A 1-A 2** and $0 < \epsilon \leq \inf_{1 \leq i \leq d} \inf_{1 \leq j \leq K^i} |c_j^i - c_{j+1}^i|^4 / 16$, where $\{c_j^i\}$ are

the quantization levels. Then, $\ell(w_k)$ converges and $\lim_{k \rightarrow \infty} d(w_k, \mathcal{Z}_\epsilon) = 0$ almost surely.

Note that the condition on ϵ ensures that the projection of C_ϵ onto the i -th coordinate is a disconnected set of K^i intervals. The proof is based on a general convergence result of Davis et al. (2020), on asymptotic behavior of stochastic approximation of differential inclusion (DI). In our particular case, the corresponding DI is $\dot{y}(t) \in -\nabla \ell(y(t)) - N_{C_\epsilon}(y(t))$ (we might notice here that this DI is also the continuous-time limit of the projected gradient method). Definitions and important results on DIs and their stochastic approximations can be found in Appendix A.1.

The proof of Theorem 1 is done in several steps (see Appendix A for complete derivations). First we prove that almost surely, the sequence of iterates (w_k) converges to C_ϵ (see Lemma 6). Then we show that an update step of ASkewSGD can be written as $w_{k+1} = w_k - \gamma_k \nabla \ell(w_k) + \gamma_k \eta_{k+1} - \gamma_k u_k$, where $\eta_{k+1} = \nabla \ell(w_k) - \widehat{\nabla} \ell(w_k)$ and u_k approximates an element of $N_{C_\epsilon}(w_k)$. We show the convergence of $\sum_{j=1}^k \gamma_j \eta_{j+1}$ in Lemma 7, and complete the proof by applying Theorem 3, which is adapted from Davis et al. (2020, Theorem 3.2).

Forward pass quantization For completeness, we finally describe the quantization of the activation function when ASkewSGD is used to train a deep NN. During the forward pass, we employ a round-to-nearest approach INT4 quantization methods for the activations, taken from Chmiel et al. (2021). We make use of Statistics Aware Weight Binning (SAWB) of (Choi et al., 2018), which finds the optimal scaling factor that minimizes the quantization error based on the statistical characteristics of activation distribution. As emphasized by (Chmiel et al., 2021; Choi et al., 2018), nonlinearities of loss and activation functions make unnecessary the use of an unbiased scalar quantizer. After scaling, we use a uniform quantization (e.g., INT4): the set of quantization values \mathcal{Q} is defined coordinate wise: $\{c_1^i, \dots, c_{K^i}^i\}$. In our experiments (see Section 4) both the weights and the activations are rescaled layerwise to fit the quantization interval (e.g., $[-2^3, 2^3]$ for INT4). The quantization values $\{c_1^i, \dots, c_{K^i}^i\}$ are the integer from the quantization interval (e.g., $\{-8, -7, \dots, 8\}$ for INT4). After quantization, both the weights and the activation are rescaled using the scaling factor calculated layerwise.

4 Experiments

We evaluate the performance of ASkewSGD with weights quantized with 1, 2, and 4 bits. While BNN performs well on some simple benchmarks, it lags significantly behind full precision NN on more demanding tasks. QNN with higher precision and quantization of the activations offers a trade-off between performance and computation efficiency. For simplicity, we refer to $[Wx/ Ay]$ as a neural architecture

with x -bit precision weights and y -bit precision activations. Details of the implementations and complementary experiments are reported in Appendix B. In all experiments, ϵ is annealed throughout the training process during successive episodes. Our experiments show that the initial value for ϵ is not critical. We use a logarithmic schedule. Given a fixed ϵ , we run the algorithm until the test error does not improve, and then reduce it by using the last iterates as the starting point for the next round. For example, in the experiments of Table 2 and Table 1, the initial value for ϵ is 1, and we reduce it as K^t with $K = 0.88$. We can set K to different values ($\frac{1}{2}$, 0.8 were tested) as long as $K < 1$.

4.1 1-bit quantization

We evaluate the performance of ASkewSGD [W1/A32] on four tasks: a convex problem, a 2D toy example and two classical image classification benchmarks.

Convex toy example We compare ASkewSGD, BinaryConnect (Courbariaux et al., 2015) and AdaSTE (Le et al., 2021) in a logistic regression problem. We generate $n = 6000$ feature vectors $\{x_k\}_{k=1}^n$ of dimension $d = 10$, drawn independently from the uniform distribution in $[-1, 1]$. We randomly choose an optimal vector w_* on the vertices of the hypercube and generate the labels as follows: $y_k \sim \text{Bernoulli}(\{1 + e^{-x_k^\top w_*}\}^{-1})$. For completeness, we study how a SGD converges with full precision to the optimal point w_* of this convex problem. All methods are trained for 25 epochs using the SGD optimizer. The learning rate is set to 1 and the gradients are computed on random batches of 1000 samples. For AdaSTE, we have used the code ¹ with the hyperparameters specified in the package for annealing. ASkewSGD performance is on par with full




Figure 2: Training losses for the logistic regression problem with batches of size 1000. BinaryConnect (green), ASkewSGD (blue), full Precision (red), AdaSTE (purple) methods. The x-axis represents the iteration index. Red points are made artificially bigger to help visualization.

precision method, while the STE variants all suffer from strong oscillations (see Shekhovtsov (2021); Shekhovtsov and Yanush (2021); Bai et al. (2018)). Figure 2 illustrates

the effects of such oscillations on the convergence. In all settings, AdaSTE converges faster than BC, but still all STE variants exhibit a larger loss compared to other methods. Additional results are reported in Appendix B.

Non-convex toy example We consider the binary classification problem on "2 moons dataset" presented in Meng et al. (2020). The training dataset consists of 2000 samples (split into 2 moon-like clusters in 2 dimensions) and 200 test samples; see Appendix B. We train a BNN with 9 neurons. In this low-dimensional environment, we can enumerate all $2^9 = 512$ possible binary configurations and select the best one(s). Our method is compared with 4 different approaches: a full precision NN, BinaryConnect (Courbariaux et al., 2015), AdaSTE (Le et al., 2021), and exhaustive search. All methods are trained for 50 epochs with logistic loss. The full precision NN is trained using the Adam optimizer (Kingma and Ba, 2014) with default hyperparameters, a learning rate of 0.1, and a batch of size 100. The BinaryConnect approach is trained using the Adam optimizer with default hyperparameters, a learning rate of 1, and a batch of size 100. The AdaSTE method is implemented using a learning rate of 1. Our method uses the same parameters as the STE method, and we set α to 4. For a single run, we plotted the training loss in Figure 3. For a fair compari-




Figure 3: Training losses for the toy non-convex problem with batches of size 100. BinaryConnect (green), ASkewSGD (blue), full Precision (red), AdaSTE (purple) methods. The x-axis represents the iteration index. Red points are made artificially bigger to help visualization

son, in Table 4 in Appendix B we report the performance averaged on 50 random experiments of the various methods on the test set (full precision reaches a 2.045 ± 0.005 loss, when exhaustive search presents a 2.1 loss, ASkewSGD reaches 2.11 ± 0.01 , AdaSTE and STE reach 2.24 ± 0.10 and 2.32 ± 0.11 respectively).

The exhaustive search shows that different configurations lead to near-optimal performance (see Figure 8 in Appendix B). Here we chose the configuration that achieves the lowest loss on the test set. ASkewSGD outperforms AdaSTE and BC.

¹<https://github.com/intellhave/AdaSTE>

Computer vision tasks In this section, we benchmark ASkewSGD with BC (Courbariaux et al., 2015; Hubara et al., 2016), Mirror Descent (Ajanthan et al., 2021), and ProxQuant (Bai et al., 2018) on classical computer vision datasets. To avoid overloading the figures, the AdaSTE results are reported separately in Appendix B. We also report performance with a standard full precision NN and a full precision NN projected onto the hypersphere. We compare the different methods using the same NN architecture. We do not add bias on any neuron. We introduce batch normalisation (without learning scale and bias parameters) after each layer. We emphasise that our method is generic and not specific to the classical ConvNet architecture. We have also obtained SOTA results for large ResNet architectures (see Table 3).

We use the standard data augmentations and normalizations for all the methods. ASkewSGD is implemented in Pytorch, and the experiments are run on a NVIDIA Tesla-P100 GPU. Standard multiclass cross-entropy loss is used for all experiments unless otherwise stated. We perform cross-validation of the hyperparameters, such as the learning rate, the trade-off between constraints α , the rate of increase of the annealing hyperparameter, and their respective schedules. The search space for tuning the hyperparameters and the final hyperparameters can be found in Appendix B. All models are fine-tuned for 100 epochs using the Adam (Kingma and Ba, 2014) optimizer with dynamics of 0.9 and 0.999, and batch of size 100.

The NN with full precision is trained with an initial learning rate of 0.08. The projected full precision NN uses a projected gradient algorithm. The same hyperparameters as the "plain" algorithm are used, except that a deterministic projection onto the hypersphere is performed for each iteration $w_{k+1} = \Pi(w_k - \gamma_k \widehat{\nabla} \ell(w_k))$. For BinaryConnect, we use the method described in Courbariaux et al. (2015). For Mirror Descent (MD), we use the code² from Ajanthan et al. (2021) and implement the version $\tanh(\cdot)$ (without annealing and with $\alpha = 0.01$ and $\mu = 100$ when training). ProxQuant was run with the parameters specified in Bai et al. (2018). Note ProxQuant does not initially quantize the fully-connected layer, and add full precision biases. For fair comparison we have tested ProxQuant with all layers binarized. The ASkewSGD method is described in Algorithm 1. Multiple values for α in $[0.1, 5]$ are considered. The precision threshold ϵ is decreased from epoch to epoch: it is set to 1 at the beginning and then exponentially annealed to $.88^t$ in the last 50 epochs, where t is the epoch. After the last step, all weights are within an interval of length $\epsilon_{\text{final}} = 0.01$ of $\{-1, +1\}$.

For ASkewSGD we apply the function $\text{sign}(\cdot)$ to our NN before evaluating it on the test set. For a fair comparison, each method was randomly initialized and independently

executed 5 times. An intensive learning rate search was also performed independently for each method. The learning rate at epochs $[20, 40]$ is divided by 2 for all methods.

Most neural networks use the inference accuracy of image classification as an evaluation metric. We first compared the training/testing accuracy with the CIFAR-10 dataset (Krizhevsky et al., 2009), which consists of 50000 training images and 10000 test images (in 10 classes). Figure 4 illustrates the distribution of the weights of the first convolutional layer (the behavior is similar for other layers) at epochs 20, 39, 55, and 99. We have also tested ASkewSGD




Figure 4: Histogram of weights during the training phase of our ASkewSGD [W1/A32] on CIFAR-10.

Table 1: Test accuracy (average over 5 random experiments) for ASkewSGD [W1/A32] at several epochs.

Epochs	ϵ	CIFAR-10	TinyImageNet
50	0.88	75.77	8.74
65	0.15	88.37	31.97
90	0.006	88.84	46.96




Figure 5: Training Loss of ASkewSGD [W1/A32] on CIFAR-10 (left) and TinyImageNet (right). The x-axis represents the batch iterations and green vertical lines correspond to epochs $[50, 65, 90]$.

[W1/A32] on the TinyImageNet dataset (Le and Yang, 2015) with a ResNet-18. TinyImageNet has 200 classes and each class has 500 (RGB) training images, 50 validation images, and 50 test images. To train ResNet-18 we follow the common practices used for training NNs: we resize the input images to 64×64 and then randomly flip them horizontally during training. During testing we center-crop them to the

²<https://github.com/kartikgupta-at-anu/md-bnn>

corresponding sizes. In Figure 5, the loss increases slightly in the final steps as the constraints become more stringent. However, this increase in training loss remains moderate and the final performance in both the training set and the test set is the best among all methods. Some test accuracies are presented in Table 1 at several epochs (identified with green lines in Figure 5) with the corresponding precision ϵ . The best test classification accuracies of the binary networks obtained with each method are listed in Table 2. For reproducibility none of the concurrent results are reported from existing papers, but each approach has been independently rerun from the available codes. Compared to other binarization algorithms, our method consistently yields better or equivalent results, while narrowing the performance gap between binary networks and floating-point counterparts on multiple datasets to an acceptable level. The performance of the projected gradient method highlights the strength of our method: we do not simply project the iterates on the nearest constraint set, but progressively push the iterates towards a smoothed version of the constraints (see Section 3), which leads to better results.

Table 2: Best Test accuracy (average and variance over 5 random experiments) after 100 training epochs.

Method	CIFAR-10	TinyImageNet
Full-precision [W32/A32]	89.46 ± 0.07	56.46 ± 0.46
BinaryConnect [W1/A32]	88.33 ± 0.29	42.35 ± 0.33
MD [W1/A32]	88.13 ± 0.25	34.89 ± 0.36
ProxQuant [W1/A32]	88.22 ± 0.28	48.79 ± 0.32
Projected gradient [W1/A32]	71.34 ± 0.46	11.78 ± 0.67
ASkewSGD [W1/A32]	88.98 ± 0.35	50.23 ± 0.37

4.2 Low-bit quantization

We consider now low-bit weight quantization and activation quantization. To fully benefit from low precision arithmetic, one should also tackle the problem of gradient quantization (Chmiel et al., 2021; Sun et al., 2020) and accumulation. We keep the last fully connected layer in full-precision, following Liu et al. (2020); Chmiel et al. (2021). We evaluate the performance of ASkewSGD [W1/A32], ASkewSGD [W2/A4], and ASkewSGD [W4/A4] on TinyImageNet and ImageNet (Russakovsky et al., 2015) datasets with a ResNet-18 network. For ASkewSGD we project NN weights onto the set of quantization values before evaluating it on the test set. For ImageNet, we keep the first convolution layer in full-precision. We use the same pre-processing (centering and data normalization) for all the methods: we resize the input images to 256×256 and then randomly crop them to 224×224 while centering them to the appropriate sizes during training. Standard multiclass cross entropy loss is used. All models are fine-tuned for 200 epochs using the Adam (Kingma and Ba, 2014) optimizer with dynamics of 0.9 and 0.999 and a batch of size 512. All methods are trained with an initial learning rate of 0.06 for TinyImageNet and 0.1

for ImageNet. The same hyperparameters are used as in the previous section for TinyImageNet. For ImageNet, the learning rate at epochs [30, 60, 90] is divided by 10 for all methods. We have run the code³ from LUQ and adapted it to TinyImageNet dataset. For a fair comparison we compute neural gradients in full precision. The results for the method Ultra-low (Sun et al., 2020) are taken from Chmiel et al. (2021).

We decided not to include the regularisation-based binarization approach (Ding et al., 2019), which addresses the activation binarization problem, in our benchmark. We have also not included in our benchmark improvements of BC methods which have been proposed in (Zhou et al., 2016; Liu et al., 2018; Bethge et al., 2020; Rastegari et al., 2016; Martinez et al., 2019); these methods are all based on the STE (Courbariaux et al., 2015) optimizer to update quantized weights. These methods have been shown to be outperformed by AdaBin (Tu et al., 2022) and ReacNet Liu et al. (2020). The latter are currently SOTA methods for energy-friendly inference on the ImageNet dataset. Note that these binary approaches still have a gap in terms of full precision performance, which needs to be addressed by modifying the NN structure (Liu et al., 2020). For ReacNet and AdaBin, we have reported the best results of Tu et al. (2022) for ResNet-18 on ImageNet.

Table 3: Best Test accuracy (single run for ImageNet due to longer training time) after 200 training epochs. * indicates the results are directly reported from existing literature.

Method	TinyImageNet	ImageNet
Full-precision [W32/A32]	56.46 ± 0.46	69.32
ReacNet [W1/A1] (2 steps)	-	65.5*
AdaBin [W1/A1] (2 steps)	-	66.4*
Ultra-low [W4/A4]	-	68.27*
LUQ [W2/A4]	54.14 ± 0.42	-
LUQ [W4/A4]	55.69 ± 0.32	68.41
ASkewSGD [W2/A4]	53.54 ± 0.28	66.45
ASkewSGD [W4/A4]	55.85 ± 0.30	68.51

ASkewSGD performs better than or on par with state of the art QNN methods and offers a shorter gap to full precision performances compared with best BNNs.

5 Conclusion

In this paper, we present ASkewSGD a novel framework for QNN training based on an annealed sequence of interval-constrained nonconvex optimization problems solved by an algorithm inspired by Muehlebach and Jordan (2021). For each of these subproblems we give theoretical guaran-

³<https://openreview.net/forum?id=clwYez4n8e8>

tees. ASkewSGD outperforms or is on par with other QNN training methods on all considered tasks.

In the current context, we estimated the carbon footprint of our experiments to be about 180 kg CO₂e (calculated using green-algorithms.org v2.1 Lannelongue et al. (2021)). This shed light on the crucial need to develop energy friendly NNs.

Acknowledgements

The authors acknowledge support of the Lagrange Mathematics and Computing Research Center.

References

- Ajanthan, T., Dokania, P. K., Hartley, R., and Torr, P. H. (2019). Proximal mean-field for neural network quantization. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 4871–4880.
- Ajanthan, T., Gupta, K., Torr, P., Hartley, R., and Dokania, P. (2021). Mirror descent view for neural network quantization. In *International Conference on Artificial Intelligence and Statistics*, pages 2809–2817. PMLR.
- Anderson, A. G. and Berg, C. P. (2018). The high-dimensional geometry of binary neural networks. In *International Conference on Learning Representations*.
- Bai, Y., Wang, Y.-X., and Liberty, E. (2018). Proxquant: Quantized neural networks via proximal operators. In *International Conference on Learning Representations*.
- Banner, R., Hubara, I., Hoffer, E., and Soudry, D. (2018). Scalable methods for 8-bit training of neural networks. *Advances in neural information processing systems*, 31.
- Beck, A. and Teboulle, M. (2000). Global optimality conditions for quadratic optimization problems with binary constraints. *SIAM Journal on Optimization*, 11(1):179–188.
- Bethge, J., Bartz, C., Yang, H., Chen, Y., and Meinel, C. (2020). Meliusnet: Can binary neural networks achieve mobilenet-level accuracy? *arXiv preprint arXiv:2001.05936*.
- Bhalgat, Y., Lee, J., Nagel, M., Blankevoort, T., and Kwak, N. (2020). Lsq+: Improving low-bit quantization through learnable offsets and better initialization. In *2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops (CVPRW)*, pages 2978–2985. IEEE Computer Society.
- Borwein, J. M. and Lewis, A. S. (2006). *Convex analysis and nonlinear optimization*, volume 3 of *CMS Books in Mathematics/Ouvrages de Mathématiques de la SMC*. Springer, second edition. Theory and examples.
- Bussieck, M. R., Drud, A. S., and Meeraus, A. (2003). Minlplib—a collection of test models for mixed-integer nonlinear programming. *INFORMS Journal on Computing*, 15(1):114–119.
- Carreira-Perpinán, M. A. and Idelbayev, Y. (2017). Model compression as constrained optimization, with application to neural nets. part ii: Quantization. *arXiv preprint arXiv:1707.04319*.
- Chen, J., Gai, Y., Yao, Z., Mahoney, M. W., and Gonzalez, J. E. (2020). A statistical framework for low-bitwidth training of deep neural networks. *Advances in Neural Information Processing Systems*, 33:883–894.
- Chmiel, B., Banner, R., Hoffer, E., Yaacov, H. B., and Soudry, D. (2021). Logarithmic unbiased quantization: Simple 4-bit training in deep learning. *arXiv preprint arXiv:2112.10769*.
- Choi, J., Chuang, P. I.-J., Wang, Z., Venkataramani, S., Srinivasan, V., and Gopalakrishnan, K. (2018). Bridging the accuracy gap for 2-bit quantized neural networks (qnn). *arXiv preprint arXiv:1807.06964*.
- Choukroun, Y., Kravchik, E., Yang, F., and Kisilev, P. (2019). Low-bit quantization of neural networks for efficient inference. In *2019 IEEE/CVF International Conference on Computer Vision Workshop (ICCVW)*, pages 3009–3018. IEEE.
- Courbariaux, M., Bengio, Y., and David, J.-P. (2015). Binaryconnect: Training deep neural networks with binary weights during propagations. In *Advances in neural information processing systems*, pages 3123–3131.
- Darabi, S., Belbahri, M., Courbariaux, M., and Nia, V. P. (2018). Bnn+: Improved binary network training.
- Davis, D., Drusvyatskiy, D., Kakade, S. M., and Lee, J. D. (2020). Stochastic subgradient method converges on tame functions. *Found Comput Math*, 20(1):119–154.
- Deng, L., Li, G., Han, S., Shi, L., and Xie, Y. (2020). Model compression and hardware acceleration for neural networks: A comprehensive survey. *Proceedings of the IEEE*, 108(4):485–532.
- Ding, R., Chin, T.-W., Liu, Z., and Marculescu, D. (2019). Regularizing activation distribution for training binarized deep networks. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 11408–11417.
- Esser, S. K., Appuswamy, R., Merolla, P., Arthur, J. V., and Modha, D. S. (2015). Backpropagation for energy-efficient neuromorphic computing. *Advances in neural information processing systems*, 28.
- Faraone, J., Fraser, N., Blott, M., and Leong, P. H. (2018). Syq: Learning symmetric quantization for efficient deep neural networks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 4300–4309.
- Fournarakis, M. and Nagel, M. (2021). In-hindsight quantization range estimation for quantized training. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 3063–3070.
- Garey, M. R. and Johnson, D. S. (1980). Computers and intractability: A guide to the theory of np -completeness. *Bull. Amer. Math. Soc*, 3:898–904.
- Gholami, A., Kim, S., Dong, Z., Yao, Z., Mahoney, M. W., and Keutzer, K. (2021). A survey of quantization methods for efficient neural network inference. *arXiv preprint arXiv:2103.13630*.
- Gill, P. E. and Wong, E. (2012). Sequential quadratic programming methods. In *Mixed integer nonlinear programming*, pages 147–224. Springer.
- Gong, R., Liu, X., Jiang, S., Li, T., Hu, P., Lin, J., Yu, F., and Yan, J. (2019). Differentiable soft quantization:

- Bridging full-precision and low-bit neural networks. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 4852–4861.
- Gu, S., Levine, S., Sutskever, I., and Mnih, A. (2016). Muprop: Unbiased backpropagation for stochastic neural networks. In *ICLR (Poster)*.
- Guo, Y. (2018). A survey on methods and theories of quantized neural networks. *arXiv preprint arXiv:1808.04752*.
- Horst, R. and Tuy, H. (2013). *Global optimization: Deterministic approaches*. Springer Science & Business Media.
- Hou, L. and Kwok, J. T. (2018). Loss-aware weight quantization of deep networks. *arXiv preprint arXiv:1802.08635*.
- Hu, Q., Wang, P., and Cheng, J. (2018). From hashing to cnns: Training binary weight networks via hashing. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 32.
- Hubara, I., Courbariaux, M., Soudry, D., El-Yaniv, R., and Bengio, Y. (2016). Binarized neural networks. *Advances in neural information processing systems*, 29.
- Jacob, B., Kligys, S., Chen, B., Zhu, M., Tang, M., Howard, A., Adam, H., and Kalenichenko, D. (2018). Quantization and training of neural networks for efficient integer-arithmetic-only inference. In *2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 2704–2713. IEEE.
- Jang, E., Gu, S., and Poole, B. (2016). Categorical reparameterization with gumbel-softmax.
- Khan, M. E. and Rue, H. (2021). The bayesian learning rule. *arXiv preprint arXiv:2107.04562*.
- Kim, M. and Smaragdis, P. (2016). Bitwise neural networks. *arXiv preprint arXiv:1601.06071*.
- Kingma, D. P. and Ba, J. (2014). Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*.
- Klenke, A. (2013). *Probability Theory: A Comprehensive Course*. Universitext. Springer London.
- Krizhevsky, A., Hinton, G., et al. (2009). Learning multiple layers of features from tiny images.
- Lanelongue, L., Grealey, J., and Inouye, M. (2021). Green algorithms: Quantifying the carbon footprint of computation. *Advanced Science*, 8(12):2100707.
- Le, H., Højer, R. K., Lin, C.-T., and Zach, C. (2021). Adaste: An adaptive straight-through estimator to train binary neural networks. *arXiv preprint arXiv:2112.02880*.
- Le, Y. and Yang, X. (2015). Tiny imagenet visual recognition challenge. *CS 231N*, 7(7):3.
- Leng, C., Dou, Z., Li, H., Zhu, S., and Jin, R. (2018). Extremely low bit neural network: Squeeze the last bit out with admm. In *Thirty-Second AAAI Conference on Artificial Intelligence*.
- Li, F., Zhang, B., and Liu, B. (2016). Ternary weight networks. *arXiv preprint arXiv:1605.04711*.
- Li, H., De, S., Xu, Z., Studer, C., Samet, H., and Goldstein, T. (2017). Training quantized nets: A deeper understanding. *Advances in Neural Information Processing Systems*, 30.
- Li, Y., Dong, X., and Wang, W. (2019). Additive powers-of-two quantization: An efficient non-uniform discretization for neural networks. *arXiv preprint arXiv:1909.13144*.
- Liang, T., Glossner, J., Wang, L., Shi, S., and Zhang, X. (2021). Pruning and quantization for deep neural network acceleration: A survey. *Neurocomputing*, 461:370–403.
- Lin, M., Ji, R., Xu, Z., Zhang, B., Wang, Y., Wu, Y., Huang, F., and Lin, C.-W. (2020). Rotated binary neural network. *Advances in neural information processing systems*, 33:7474–7485.
- Liu, Z., Shen, Z., Savvides, M., and Cheng, K.-T. (2020). Reactnet: Towards precise binary neural network with generalized activation functions. In *European Conference on Computer Vision*, pages 143–159. Springer.
- Liu, Z., Wu, B., Luo, W., Yang, X., Liu, W., and Cheng, K.-T. (2018). Bi-real net: Enhancing the performance of 1-bit cnns with improved representational capability and advanced training algorithm. In *Proceedings of the European conference on computer vision (ECCV)*, pages 722–737.
- Maddison, C., Mnih, A., and Teh, Y. (2017). The concrete distribution: A continuous relaxation of discrete random variables. In *Proceedings of the international conference on learning Representations*. International Conference on Learning Representations.
- Martinez, B., Yang, J., Bulat, A., and Tzimiropoulos, G. (2019). Training binary neural networks with real-to-binary convolutions. In *International Conference on Learning Representations*.
- McDonnell, M. D. (2018). Training wide residual networks for deployment using a single bit for each weight. *arXiv preprint arXiv:1802.08530*.
- Meng, X., Bachmann, R., and Khan, M. E. (2020). Training binary neural networks using the bayesian learning rule. In *International conference on machine learning*, pages 6852–6861. PMLR.
- Mishra, A. and Marr, D. (2017). Apprentice: Using knowledge distillation techniques to improve low-precision network accuracy. *arXiv preprint arXiv:1711.05852*.
- Mishra, A., Nurvitadhi, E., Cook, J. J., and Marr, D. (2017). Wrpn: Wide reduced-precision networks. *arXiv preprint arXiv:1709.01134*.
- Mitchell, J. E., Pardalos, P. M., and Resende, M. G. (1998). Interior point methods for combinatorial optimization. In *Handbook of combinatorial optimization*, pages 189–297. Springer.

- Mnih, A. and Gregor, K. (2014). Neural variational inference and learning in belief networks. In *International Conference on Machine Learning*, pages 1791–1799. PMLR.
- Moons, B., Goetschalckx, K., Van Berckelaer, N., and Verhelst, M. (2017). Minimum energy quantized neural networks. In *2017 51st Asilomar Conference on Signals, Systems, and Computers*, pages 1921–1925. IEEE.
- Muehlebach, M. and Jordan, M. I. (2021). On constraints in first-order optimization: A view from non-smooth dynamical systems. *arXiv preprint arXiv:2107.08225*.
- Murray, W. and Ng, K.-M. (2010). An algorithm for nonlinear optimization problems with binary variables. *Computational optimization and applications*, 47(2):257–288.
- Nagel, M., Fournarakis, M., Amjad, R. A., Bondarenko, Y., van Baalen, M., and Blankevoort, T. (2021). A white paper on neural network quantization. *arXiv preprint arXiv:2106.08295*.
- Peters, J. W. and Welling, M. (2018). Probabilistic binary neural networks. *arXiv preprint arXiv:1809.03368*.
- Polino, A., Pascanu, R., and Alistarh, D. (2018). Model compression via distillation and quantization. *arXiv preprint arXiv:1802.05668*.
- Qin, H., Gong, R., Liu, X., Bai, X., Song, J., and Sebe, N. (2020). Binary neural networks: A survey. *Pattern Recognition*, 105:107281.
- Raiko, P., Berglund, M., Alain, G., and Dinh, L. (2015). Techniques for learning binary stochastic feedforward neural networks. In *International Conference on Learning Representations*, pages 1–10.
- Rastegari, M., Ordonez, V., Redmon, J., and Farhadi, A. (2016). XNOR-Net: Imagenet classification using binary convolutional neural networks. In *European Conference on Computer Vision*, pages 525–542. Springer.
- Roth, W., Schindler, G., Fröning, H., and Pernkopf, F. (2019). Training discrete-valued neural networks with sign activations using weight distributions. In *Joint European Conference on Machine Learning and Knowledge Discovery in Databases*, pages 382–398. Springer.
- Russakovsky, O., Deng, J., Su, H., Krause, J., Satheesh, S., Ma, S., Huang, Z., Karpathy, A., Khosla, A., Bernstein, M., Berg, A. C., and Fei-Fei, L. (2015). ImageNet Large Scale Visual Recognition Challenge. *International Journal of Computer Vision (IJCV)*, 115(3):211–252.
- Sakr, C. and Shanbhag, N. (2018). Per-tensor fixed-point quantization of the back-propagation algorithm. *arXiv preprint arXiv:1812.11732*.
- Sard, A. (1942). The measure of the critical values of differentiable maps. *Bulletin of the American Mathematical Society*, 48(12):883 – 890.
- Shekhovtsov, A. (2021). Bias-variance tradeoffs in single-sample binary gradient estimators. In *DAGM German Conference on Pattern Recognition*, pages 127–141. Springer.
- Shekhovtsov, A. and Yanush, V. (2021). Reintroducing straight-through estimators as principled methods for stochastic binary networks. In *DAGM German Conference on Pattern Recognition*, pages 111–126. Springer.
- Sun, X., Wang, N., Chen, C.-Y., Ni, J., Agrawal, A., Cui, X., Venkataramani, S., El Maghraoui, K., Srinivasan, V. V., and Gopalakrishnan, K. (2020). Ultra-low precision 4-bit training of deep neural networks. *Advances in Neural Information Processing Systems*, 33:1796–1807.
- Tu, Z., Chen, X., Ren, P., and Wang, Y. (2022). Adabin: Improving binary neural networks with adaptive binary sets. *arXiv preprint arXiv:2208.08084*.
- Wang, N., Choi, J., Brand, D., Chen, C.-Y., and Gopalakrishnan, K. (2018). Training deep neural networks with 8-bit floating point numbers. *Advances in neural information processing systems*, 31.
- Wang, Z., Lu, J., Tao, C., Zhou, J., and Tian, Q. (2019). Learning channel-wise interactions for binary convolutional neural networks. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 568–577.
- Wiedemann, S., Mehari, T., Kepp, K., and Samek, W. (2020). Dithered backprop: A sparse and quantized backpropagation algorithm for more efficient deep neural network training. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops*, pages 720–721.
- Yang, T.-J., Chen, Y.-H., and Sze, V. (2017). Designing energy-efficient convolutional neural networks using energy-aware pruning. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 5687–5695.
- Zagoruyko, S. and Komodakis, N. (2016). Wide residual networks. *arXiv preprint arXiv:1605.07146*.
- Zhang, D., Yang, J., Ye, D., and Hua, G. (2018). Lq-nets: Learned quantization for highly accurate and compact deep neural networks. In *Proceedings of the European conference on computer vision (ECCV)*, pages 365–382.
- Zhang, Z., Shao, W., Gu, J., Wang, X., and Luo, P. (2021). Differentiable dynamic quantization with mixed precision and adaptive resolution. In *International Conference on Machine Learning*, pages 12546–12556. PMLR.
- Zhao, W., Ma, T., Gong, X., Zhang, B., and Doermann, D. (2020). A review of recent advances of binary neural networks for edge computing. *IEEE Journal on Miniaturization for Air and Space Systems*.
- Zhou, A., Yao, A., Guo, Y., Xu, L., and Chen, Y. (2017). Incremental network quantization: Towards loss-

less cnns with low-precision weights. *arXiv preprint arXiv:1702.03044*.

Zhou, S., Wu, Y., Ni, Z., Zhou, X., Wen, H., and Zou, Y. (2016). Dorefa-net: Training low bitwidth convolutional neural networks with low bitwidth gradients. *arXiv preprint arXiv:1606.06160*.

Zhou, Y., Moosavi-Dezfooli, S.-M., Cheung, N.-M., and Frossard, P. (2018). Adaptive quantization for deep neural network. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 32.

Zhu, C., Han, S., Mao, H., and Dally, W. J. (2016). Trained ternary quantization. *arXiv preprint arXiv:1612.01064*.

A Proofs of Section 3

A.1 Preliminaries

Absolutely continuous curves. We say that a curve $y : \mathbb{R}_+ \rightarrow \mathbb{R}^d$ is absolutely continuous (a.c.) if there is a curve $z : \mathbb{R}_+ \rightarrow \mathbb{R}^d$, locally Lebesgue integrable, such that for every $t \geq 0$,

$$y(t) - y(0) = \int_0^t z(u) du.$$

In this case, it holds that for almost every $t \geq 0$, y is differentiable and $\dot{y}(t) = z(t)$.

Tangent and normal cones. Let $C \subset \mathbb{R}^d$ be a closed set. For $w \in C$, the tangent cone of C to w , denoted by $T_C(w)$, is the set of vectors $v \in \mathbb{R}^d$ for which there exist $t_k \downarrow 0$ and $w_k \rightarrow w$, $w_k \in C$, such that $(w_k - w)/t_k \rightarrow v$. The normal cone of C at w , denoted $N_C(w)$, is the set of vectors $u \in \mathbb{R}^d$ such that for any $v \in T_C(w)$, $u^\top v \leq 0$. If $w \notin C$, then by convention $T_C(w), N_C(w) = \emptyset$.

The Mangasarian-Fromovitz constraint qualification (MFCQ) condition. Consider the case where $C = \{w \in \mathbb{R}^d : g(w) \geq 0\}$, for a smooth function $g : \mathbb{R}^d \rightarrow \mathbb{R}^{n_g}$. Denote $I(w) = \{i \in \{1, \dots, n_g\}, g_i(w) \leq 0\}$ as the set of active constraints. We say that the MFCQ condition holds at $w \in \mathbb{R}^d$ if there exists $v \in \mathbb{R}^d$ such that $\nabla g_i(w)^\top v \geq 0$ for all $i \in I(w)$. If the MFCQ condition holds at $w \in C$, then we can write down $T_C(w) = \{v \in \mathbb{R}^d, \nabla g_i(w)^\top v \geq 0, \text{ for all } i \in I(w)\}$ and $N_C(w) = \{-\sum_{i=1}^{n_g} \lambda_i \nabla g_i(w), \lambda_i \in \mathbb{R}_+ \text{ and } \lambda_i = 0 \text{ if } i \notin I(w)\}$ (see, e.g., (Borwein and Lewis, 2006, Section 7.2)). We might notice here, that in the context of Theorem 1 the MFCQ condition holds at every $w \in C_\epsilon$.

Differential inclusion. Consider a closed set $C \subset \mathbb{R}^d$ and $\ell : \mathbb{R}^d \rightarrow \mathbb{R}$ a smooth function. An essential ingredient of our proof will be the following differential inclusion (DI):

$$\dot{y}(t) \in -\nabla \ell(y(t)) - N_C(y(t)). \quad (4)$$

We say that an a.c. curve $y : \mathbb{R}_+ \rightarrow C$ is a solution to this DI if the inclusion holds for almost every $t \geq 0$. We say that ℓ is a Lyapunov function for the set $\mathcal{Z} := \{w \in \mathbb{R}^d : 0 \in -\nabla \ell(w) - N_C(w)\}$ if for any such curve:

$$\text{for all } t > 0, \quad \ell(y(t)) \leq \ell(y(0)),$$

with strict inequality as soon as $y(0) \notin \mathcal{Z}$. We have the following lemma.

Lemma 2. *Assume that MFCQ holds at every $w \in C$. Then ℓ is a Lyapunov function for the DI (4) and the set \mathcal{Z} .*

Proof. The assumption that MFCQ holds at every $w \in C$ implies that C is Clarke regular (i.e. if $(w_k, u_k) \rightarrow (w, u) \in C \times \mathbb{R}^d$ with $(w_k, u_k) \in C \times N_C(w_k)$, then $u \in N_C(w)$). As shown in (Sections 5 and 6 Davis et al. (2020)), this implies that for almost every $t \geq 0$ and every $v \in N_C(y(t))$, $\dot{y}(t)^\top v = 0$. Therefore, for almost every $t \geq 0$,

$$\begin{aligned} \frac{d}{dt} \ell(y(t)) &= \nabla \ell(y(t))^\top \dot{y}(t) \\ &\in -\|\dot{y}(t)\|^2 - \dot{y}(t)^\top v(t) = -\|\dot{y}(t)\|^2, \end{aligned}$$

where $v(t) = \nabla \ell(y(t)) - \dot{y}(t) \in -N_C(y(t))$. This shows that $y(t) - y(0) = -\int_0^t \|\dot{y}(u)\|^2 du$, which, by closedness of \mathcal{Z} , implies our statement. \square

In Section 3 the set of interest will be C_ϵ . It can be easily seen that under the assumptions of Theorem 1 the MFCQ condition is satisfied at every $w \in C_\epsilon$. Thus, Lemma 2 implies that, in this context, ℓ is a Lyapunov function for the DI: $\dot{y}(t) \in -\ell(y(t)) - N_{C_\epsilon}(y(t))$.

Discrete approximations of differential inclusions. The idea of our proof is to apply the results of Davis et al. (2020) on the stochastic approximation of differential inclusions to our setting. To this end, we consider an \mathbb{R}^d -valued sequence (y_k) constructed as follows:

$$y_{k+1} = y_k - \gamma_k \nabla \ell(y_k) + \gamma_k \eta_{k+1} - \gamma_k u_k,$$

where (γ_k) is a sequence of positive step-sizes and $(\eta_k), (u_k)$ are some \mathbb{R}^d -valued sequences. Here, u_k represent some approximation of an element of $N_C(y_k)$, and η_{k+1} some (stochastic or deterministic) perturbation. Therefore, (y_k) might be seen as an Euler-like discretization of the DI (4).

The following proposition follows from a general result of (Davis et al., 2020, Theorem 3.2). We state it, applied to our particular case.

Theorem 3. *Assume that:*

1. *The sequence (γ_k) satisfies $\sum_{j=0}^{+\infty} \gamma_j = +\infty$ and $\sum_{j=0}^{+\infty} \gamma_j^2 < +\infty$.*
2. *The sequence $\left(\sum_{j=0}^n \gamma_j \eta_{j+1}\right)$ converges.*
3. *The sequence (y_k, u_k) is bounded.*
4. *If y_{k_j} is a subsequence such that $y_{k_j} \rightarrow y_\infty$, then $y_\infty \in C$ and the distance between $-N_C(y_\infty) - \nabla \ell(y_\infty)$ and $-1/n \sum_{j=1}^n \{\nabla \ell(y_{k_j}) + u_{k_j}\}$ goes to zero.*
5. *ℓ is a Lyapunov function for the DI (4).*
6. *The set $\ell(\mathcal{Z})$ is of empty interior.*

Then, $\ell(y_k)$ converges and $\limsup_{k \rightarrow +\infty} d(y_k, \mathcal{Z}) = 0$.

Proof. Apply (Davis et al., 2020, Theorem 3.2), with $G = -\nabla \ell - N_C$ and $\phi = \ell$. □

Lemma 4. *We can replace the 4-th assumption in Theorem 3 by the following assumption: if (y_∞, u_∞) is a cluster point of (y_k, u_k) , then $u_\infty \in N_C(y_\infty)$.*

Proof. If (y_{k_j}) is a subsequence such that $y_{k_j} \rightarrow y_\infty$, then $1/n \sum_{j=1}^n -\nabla \ell(y_{k_j}) \rightarrow_{n \rightarrow \infty} -\nabla \ell(y_\infty)$. Furthermore, for any $m \geq 0$, we can write:

$$\frac{1}{n} \sum_{j=1}^n u_{k_j} = \frac{1}{n} \sum_{j=1}^m u_{k_j} + \frac{n-m}{n} \left(\frac{1}{n-m} \sum_{j=m+1}^n u_{k_j} \right).$$

By the Caratheodory theorem, we can write $1/(n-m) \sum_{j=m+1}^n u_{k_j} = \sum_{i=1}^{d+1} \lambda_{m,n,i} u_{m,n,i}$, where $\lambda_{m,n,i} \geq 0$, $\sum_{i=1}^{d+1} \lambda_{m,n,i} = 1$ and $u_{m,n,i} \in \{u_{k_{m+1}}, \dots, u_{k_n}\}$. Denote $\mathcal{C} \subset N_C(y_\infty)$ the set of cluster points of the sequence u_{k_j} . Since the sequence (u_k) is bounded, for each $i \in \{1, \dots, d+1\}$, we can extract a convergent sequence from $(\lambda_{m,n,i}, u_{m,n,i})$ that converges to $(\lambda_m(i), u_m(i))$, with $u_m(i) \in \mathcal{C} \cup \bigcup_{j=m+1}^{+\infty} \{u_{k_j}\}$. Thus, $1/n \sum_{j=1}^n u_{k_j} \rightarrow \sum_{i=1}^{d+1} \lambda_m(i) u_m(i)$. As a consequence, we can write:

$$\lim_{n \rightarrow +\infty} \frac{1}{n} \sum_{j=1}^n u_{k_j} = \lim_{m \rightarrow \infty} \sum_{i=1}^{d+1} \lambda_m(i) u_m(i).$$

For each $i \in \{1, \dots, d+1\}$, the sequences $(\lambda_m(i))_{m \geq 0}, (u_m(i))_{m \geq 0}$ are bounded. Therefore, up to an extraction of a subsequence, we can assume that they converge to some $\lambda(i), u(i)$. Notice that $u(i) \in \mathcal{C} \subset N_C(y_\infty)$. Therefore, $1/n \sum_{j=1}^n u_{k_j}$ converges to a convex combination of elements of $N_C(y_\infty)$. By convexity of $N_C(y_\infty)$ this implies that $1/n \sum_{j=1}^n u_{k_j}$ converges to an element of $N_C(y_\infty)$. □

The following lemma provides a condition under which $\ell(\mathcal{Z})$ has an empty interior.

Lemma 5. *Assume that $\ell : \mathbb{R}^d \rightarrow \mathbb{R}$ is d -times continuously differentiable and that $C = [a_1, b_1] \times \dots \times [a_d, b_d]$, where for $1 \leq i \leq d$, a_i, b_i are some real numbers. Consider $\mathcal{Z} = \{y \in \mathbb{R}^d : 0 \in -\nabla \ell(y) - N_C(y)\}$. It holds that $\ell(\mathcal{Z})$ is of empty interior.*

Proof. Denote \mathring{C} as the interior of C . The fact that $\ell(\mathcal{Z} \cap \mathring{C})$ has an empty interior is a consequence of Sard's theorem and the fact that ℓ is d -times differentiable (see Sard (1942)). We now show that the image of ℓ of any m -dimensional boundary of C intersected by \mathcal{Z} also has an empty interior. Consider $m > 0$, and fix $m-d$ coordinates of C as c_{m+1}, \dots, c_d , where c_i is equal to a_i or b_i , and denote $C_m = (a_1, b_1) \times (a_2, b_2) \times \dots \times (a_m, b_m) \times \{c_{m+1}\} \times \dots \times \{c_d\}$. Note that if $y \in \mathcal{Z} \cap C_m$, then the m first coordinates of $\nabla \ell(y)$ are zero. Thus, if we call ℓ_m the restriction of ℓ to C_m , then $\ell_m : C_m \rightarrow \mathbb{R}$ is d times differentiable and $\mathcal{Z} \cap C_m$ is included in its set of critical points. Applying Sard's theorem to ℓ_m , we obtain that $\ell(\mathcal{Z} \cap C_m)$ has an empty interior. Since C can be written as a union of these C_m , this completes the proof. □

A.2 A proof of Theorem 1

First we need to prove that the cluster point of the iterates w_∞ belongs to the constraints set C_ϵ .

Lemma 6. *Under assumptions of Theorem 1 it holds that $\limsup_{k \rightarrow \infty} d(w_k, C_\epsilon) = 0$ almost surely.*

Sketch of proof. The detailed proof is given in the following section Appendix A.3. The main idea is that for any $i \in \{1, \dots, d\}$, if $\psi_\epsilon^i(w_k^i) < 0$ (i.e. w_k^i is outside of the constraints), then w_k^i is constantly pushed to the closest interval. Thus, the non-convergence might happen if and only if w_k^i visits one of the interval infinitely often. However, due to the fact, that γ_k decreases to zero and that $\nabla \ell_j$ is bounded, this implies, for k large enough, that w_k^i will never leave the ‘‘region of attraction’’ of this interval (it will be kept at a distance of order γ_k to this interval) and thus converge to it. \square

Proof of Theorem 1. Our goal is to apply Theorem 3 and, hence, verify its assumptions. By a standard Martingale argument it holds that the sequence $\sum_{j=0}^k \gamma_j \eta_{j+1}$, almost surely, converges to a finite random variable (a short proof of this result is given in Appendix A.3). Consider a realization for which $\sum_{j=0}^\infty \gamma_j \eta_{j+1} < \infty$. Let (w_∞, u_∞) be a cluster point of the sequence (w_k, u_k) and let $(k_j)_{j \geq 0}$ be a subsequence such that $\lim_{j \rightarrow +\infty} (w_{k_j}, u_{k_j}) = (w_\infty, u_\infty)$. Lemma 6 shows that $w_\infty \in C_\epsilon$. Since $\sup_{k \geq k_{0,\epsilon}} |\lambda_k^i| < +\infty$, we can extract a subsequence from k_j , and assume that $\lambda_{k_j} \rightarrow \lambda$. Thus, $u_\infty^i = -\lambda^i \psi'_\epsilon(w_\infty^i)$. Since all of the $\lambda_{k_j}^i$ are positive, it holds that $\lambda^i \geq 0$. Moreover, notice that if $\psi_\epsilon(w_\infty^i) > 0$, then, for j large enough, $\psi_\epsilon(w_{k_j}^i) > 0$ and, therefore, $\lambda_{k_j}^i = 0$. Hence, for $i \notin I(w_\infty)$, $\lambda^i = 0$. This shows $u_\infty \in N_{C_\epsilon}(w_\infty)$. As shown in Lemma 2, ℓ is a Lyapunov function for the DI: $\dot{y}(t) \in -\nabla \ell(y(t)) - N_{C_\epsilon}(y(t))$. In Lemma 5 we show that $\ell(\mathcal{Z})$ is of empty interior. Thus, with the help of Lemmas 6 and 7, the assumptions of Theorem 3 are satisfied, which concludes the proof. \square

A.3 A martingale result and proof of Lemma 6

We first establish a result on the convergence of the weighted sequence of perturbations.

Lemma 7. *Assume A1-A2. Then, almost surely, $\sum_{j=0}^k \gamma_j \eta_{j+1}$ converges.*

Proof. Denote by \mathcal{F}_k the filtration generated by $\{w_1, \dots, w_k\}$. It holds that $\mathbb{E}[\widehat{\nabla} \ell(w_k) | \mathcal{F}_k] = \nabla \ell(w_k)$. Furthermore, almost surely, $\mathbb{E}[\|\eta_{k+1}\|^2 | \mathcal{F}_k] \leq 2\mathbb{E}[\|\widehat{\nabla} \ell(w_k)\|^2 | \mathcal{F}_k] + 2\|\nabla \ell(w_k)\|^2 < 4M_\ell$, where $M_\ell = \sup_{1 \leq j \leq N} M_{\ell_j}$. Thus, for $i \in \{1, \dots, d\}$, $\sum_{j=0}^k \gamma_j \eta_{j+1}^i$ is a martingale with an almost surely bounded square variation (since $\sum_{j=0}^{+\infty} \gamma_j^2 < +\infty$). The proof is concluded by applying (Klenke, 2013, Theorem 11.14). \square

In all the sequel, it is implicitly assumed that ϵ was chosen small enough to satisfy the assumption of Theorem 1. Denote by $k_{0,\epsilon}$ the smallest integer after which we do not perform the clipping step in Algorithm 1.

$$k_{0,\epsilon} := \inf \{k \geq 0 : \text{for } m \geq k, s_{\epsilon,\alpha}^c(\widehat{\nabla} \ell(w_m), w_m) = s_{\epsilon,\alpha}(\widehat{\nabla} \ell(w_m), w_m)\}.$$

Since $\limsup d(w_k, C_\epsilon) = 0$, it holds that $\liminf \psi_\epsilon(w_k^i) \geq 0$ and, therefore, $k_{0,\epsilon}$ is almost surely finite. Thus, for $k \geq k_{0,\epsilon}$, $v_k^i = [s_{\epsilon,\alpha}(\widehat{\nabla} \ell(w_k), w_k)]^i$, which implies:

$$v_k^i = -\widehat{\nabla}_i \ell(w_k) + \lambda_k^i \psi'_\epsilon(w_k^i),$$

with $\lambda_k^i = 0$ if $\psi_\epsilon(w_k^i) > 0$ and $\lambda_k^i = (v_k^i + \widehat{\nabla}_i \ell(w_k)) / \psi'_\epsilon(w_k^i)$ otherwise. Notice that since the sequences (v_k) , (w_k) are almost surely bounded, $\sup_{k \geq k_{0,\epsilon}} |\lambda_k^i|$ is almost surely finite.

Lemma 8. *Assume A1-A2. For $i \in \{1, \dots, d\}$, and for $k \geq k_{0,\epsilon}$, $\lambda_k^i \geq 0$.*

Proof. First, notice that if $\psi_\epsilon(w_k^i) > 0$, then $\lambda_k^i = 0$ by construction. Consider now the case where $\psi_\epsilon(w_k^i) \leq 0$. If $-\widehat{\nabla}_i \ell(w_k) \psi'_\epsilon(w_k^i) \geq -\alpha \psi_\epsilon(w_k^i)$, then $v_k^i = -\widehat{\nabla}_i \ell(w_k)$ and, since for $k \geq k_{0,\epsilon}$, $\psi'_\epsilon(w_k^i) \neq 0$, this implies $\lambda_k^i = 0$. Otherwise, $v_k^i = -\alpha \psi_\epsilon(w_k^i) / \psi'_\epsilon(w_k^i)$ and $0 < -\alpha \psi_\epsilon(w_k^i) + \widehat{\nabla}_i \ell(w_k) \psi'_\epsilon(w_k^i)$. Dividing the last inequality by $\{\psi'_\epsilon(w_k^i)\}^2$, we obtain $0 < (-\alpha \psi_\epsilon(w_k^i) + \widehat{\nabla}_i \ell(w_k) \psi'_\epsilon(w_k^i)) / \{\psi'_\epsilon(w_k^i)\}^2 = (v_k^i + \widehat{\nabla}_i \ell(w_k)) / \psi'_\epsilon(w_k^i) = \lambda_k^i$. \square

The rest of this section is devoted to the proof of Lemma 6.

Denote $M = \max(M_{\epsilon, c}, \sup_{1 \leq j \leq N} M_{\ell_j})$ and notice that for any $k \geq 0$ and $i \in \{1, \dots, d\}$, $\|\widehat{\nabla} \ell(w_k)\| \leq M$ and $|v_k^i| \leq M$. Therefore, $|w_{k+1}^i - w_k^i| \leq \gamma_k M$. The lemma will be proved by the following claims.

Claim 1. For $i \in \{1, \dots, d\}$, and for $2 \leq j \leq K_i - 1$ if the set $[(c_j^i + c_{j-1}^i)/2, (c_j^i + c_{j+1}^i)/2]$ is visited by w_k^i infinitely often, then there is k_0 such that for all $k > k_0$, $w_k^i \in [(c_j^i + c_{j-1}^i)/2, (c_j^i + c_{j+1}^i)/2]$.

Indeed, fix such a j and denote $[c_-, c_+]$ the set $C_\epsilon^i \cap [(c_j^i + c_{j-1}^i)/2, (c_j^i + c_{j+1}^i)/2]$, where C_ϵ^i is the projection of C_ϵ onto the i -th coordinate. Define $k_0 = \sup\{k : \gamma_k M \geq \max(c_- - (c_j^i + c_{j-1}^i)/2, (c_j^i + c_{j+1}^i)/2 - c_+)\}$. Consider $k \geq k_0$, if $(c_j^i + c_{j-1}^i)/2 \leq w_k^i \leq c_-$ (we are on the left side of the interval), then the iterate is pushed to the right and $w_k^i \leq w_{k+1}^i$. Furthermore, by definition of k_0 , it holds that $w_{k+1}^i \leq c_- + \gamma_k M \leq (c_j^i + c_{j+1}^i)/2$. This implies, that in this case w_{k+1}^i stays in $[(c_j^i + c_{j-1}^i)/2, (c_j^i + c_{j+1}^i)/2]$. Otherwise, if $c_+ \leq w_k^i < (c_j^i + c_{j+1}^i)/2$ (we are on the right side of the interval), then, we are pushed to the left, and, by a similar reasoning, $w_{k+1}^i \in [(c_j^i + c_{j-1}^i)/2, (c_j^i + c_{j+1}^i)/2]$. Finally, if $w_k^i \in [c_-, c_+]$, then by the way k_0 was defined we obtain that $w_{k+1}^i \in [(c_j^i + c_{j-1}^i)/2, (c_j^i + c_{j+1}^i)/2]$. Thus, we have shown that for $k \geq k_0$, if w_k^i is in $[(c_j^i + c_{j-1}^i)/2, (c_j^i + c_{j+1}^i)/2]$, then for all $k' \geq k$, the same will be true for $w_{k'}^i$, which completes the proof of the claim.

The proof of the following two claims is similar to the one of Claim 1.

Claim 2. For $i \in \{1, \dots, d\}$, if the set $(-\infty, (c_1^i + c_2^i)/2)$ is visited by w_k^i infinitely often, then there is k_0 such that for all $k > k_0$, $w_k^i \in (-\infty, (c_1^i + c_2^i)/2)$.

Claim 3. For $i \in \{1, \dots, d\}$, if the set $[(c_{K_i-1}^i + c_{K_i}^i)/2, +\infty)$ is visited infinitely often, then there is k_0 such that for all $k > k_0$, $w_k^i \in [(c_{K_i-1}^i + c_{K_i}^i)/2, +\infty)$.

In the following, without loss of generality, we will assume that we are in the context of the first claim and that there is k_0 , such that for all $k \geq k_0$, $w_k^i \in [(c_j^i + c_{j-1}^i)/2, (c_j^i + c_{j+1}^i)/2]$ (the two other cases can be treated in the exact same manner).

Denote, as previously, $[c_-, c_+]$ the set $C_\epsilon^i \cap [(c_j^i + c_{j-1}^i)/2, (c_j^i + c_{j+1}^i)/2]$, where C_ϵ^i is the projection of C_ϵ onto the i -th coordinate.

Claim 4. There is k_0 , such that if there are two index $m_+ \geq m_- > k_0$ such that $w_{m_-} < c_- < c_+ < w_{m_+}$, then there is m , satisfying $m_- \leq m \leq m_+$, such that $w_m^i \in [c_-, c_+]$.

Indeed, define $k_0 = \sup\{k : \gamma_k M \geq c_+ - c_-\}$. Let m_-, m_+ be as in the claim and consider $m = \inf\{k \geq m_- : w_k^i \geq c_-\}$. It holds that $w_{m-1}^i < c_- \leq w_m^i \leq w_{m-1}^i + \gamma_k M$. Since $m \geq k_0$, this implies that $w_m^i \leq c_- + \gamma_k M \leq c_+$, which proves the claim.

Claim 5. There is k_0 , such that if there are two index $m_- \geq m_+ > k_0$, such that $w_{m_-} < c_- < c_+ < w_{m_+}$, then there is m , satisfying $m_+ \leq m \leq m_-$, such that $w_m^i \in [c_-, c_+]$. The proof is the identical to the one of the previous claim.

From the fourth and fifth claims, there are only three possible behaviors of w_k^i . Either, w_k^i visits $[c_-, c_+]$ infinitely often (this will be treated by the sixth claim), or for k large enough, w_k^i stays at the left of $[c_-, c_+]$ (this will be treated by the seventh claim), or it stays at the right of $[c_-, c_+]$ (this will be treated by the eighth claim).

Claim 6. If w_k^i visits $[c_-, c_+]$ infinitely often, then $\limsup w_k^i \leq c_+$ and $\liminf w_k^i \geq c_-$.

Notice that if $w_k^i > c_+$, then $w_{k+1}^i \leq w_k^i$, and if $w_k^i \leq c_+$ and $w_{k+1}^i \leq c_+ + \gamma_k M$. Thus, if k is such that $w_k^i \in [c_-, c_+]$, then $\sup_{k_1 \geq k} w_{k_1}^i \leq c_+ + \gamma_k M$. Letting k tend to infinity, proves first part of the claim. Similarly, if k is such that $w_k^i \in [c_-, c_+]$, then $\inf_{k_1 \geq k} w_{k_1}^i \geq c_- - \gamma_k M$. Letting k tend to infinity proves the second part of the claim.

Claim 7. If for all k large enough, $w_k^i > c_+$, then $w_k^i \rightarrow c_+$.

Indeed, in this case, for k large enough, the sequence w_k^i is decreasing and thus has a limit. Denote this limit w_+ and assume that $w_+ \neq c_+$, then for k large enough, it holds that $w_{k+1}^i = w_k^i + \gamma_k v_k^i \leq w_k^i - \gamma_k M_+$, where $M_+ = \inf\{\min(M_{\epsilon, c}, \alpha|\psi_\epsilon(w)|/|\psi'_\epsilon(w)|) : w \in [w_+, (c_j + c_{j+1})/2]\} > 0$. Thus, for any m , it holds that $w_{k+m+1}^i \leq w_k^i - M_+ \sum_{i=0}^m \gamma_{k+i}$. Since $\sum_{j=0}^{+\infty} \gamma_j = +\infty$, this shows that this case is impossible. Hence, $w_k^i \rightarrow c_+$.

Claim 8. If for all k large enough, $w_k^i < c_-$, then $w_k^i \rightarrow c_-$.

Similarly, to the previous claim, for k large enough the sequence w_k^i is increasing and thus has a limit. If $w_- \neq c_-$, then for k large enough and $m \geq 0$, it holds that $w_{k+m+1}^i \geq w_k^i + M_- \sum_{i=0}^m \gamma_{k+i}$, where $M_- = \inf\{\min(M_{\epsilon, c}, \alpha|\psi_\epsilon(w)|/|\psi'_\epsilon(w)|) : w \in [w_-, (c_j + c_{j+1})/2]\} > 0$. Thus, for any m , it holds that $w_{k+m+1}^i \geq w_k^i + M_- \sum_{i=0}^m \gamma_{k+i}$. Since $\sum_{j=0}^{+\infty} \gamma_j = +\infty$, this shows that this case is impossible. Hence, $w_k^i \rightarrow c_-$.

Table 4: Logistic Loss on test samples (average over 50 random experiments) for the binary classification problem in dimension $d = 512$ after 100 steps.

Method	Loss ($\times 10^{-3}$)
Full-precision NN [W32/A32]	2.045 \pm 0.005
BinaryConnect [W1/A32]	2.32 \pm 0.11
AdaSTE [W1/A32]	2.24 \pm 0.10
ASkewSGD [W1/A32]	2.11 \pm 0.01
Exhaustive search [W1/A32]	2.1

$w \in ((c_{j-1} + c_j)/2, w_-] \} > 0$. Since $\sum_{j=0}^{+\infty} \gamma_j = +\infty$, this implies that $w_- \neq c_-$ is impossible. Hence, $w_k^i \rightarrow c_-$.

These claims show that for every $i \in \{1, \dots, d\}$, $\liminf \psi_\epsilon(w_k^i) \geq 0$. Therefore, $\limsup d(w_k, C_\epsilon) = 0$.

B Numerical results

In this section, we give more details about our experiments, and present results on new tasks.

B.1 Toy convex example

We give more results about the toy example detailed in Section 4. We only compare ASkewSGD and BinaryConnect Courbariaux et al. (2015) in a logistic regression problem, but we test several settings to highlight the strengths of ASkewSGD : all methods are trained for a longer time (50 epochs) using the SGD optimizer, the learning rate is set to 1, and gradients are calculated on random batches of 100 or 1000 samples. Note the rest of the experimental setting is identical: we generate $n = 6000$ feature vectors $\{x_k\}_{k=1}^n$ in dimension $d = 10$ drawn independently from the uniform distribution in $[-1, 1]$. We randomly choose an optimal vector w_* on the vertices of the hypercube and generate the labels as follows: $y_k \sim \text{Bernoulli}(\{1 + e^{-x_k^\top w_*}\}^{-1})$. For completeness, we study how a full precision SGD converges to the optimal point w_* of this convex problem. The same conclusions can be drawn: ASkewSGD is very close to the full precision method



Figure 6: Training losses for the logistic regression problem with batches of size 1000 (up panel) and 100 (down panel). BinaryConnect - green - ASkewSGD - blue - full Precision methods - red -. The x-axis represents the iteration index.

while STE method suffers from oscillations. Note however that decreasing the batch size seems to have a beneficial effect for STE, the larger variance helps to reduce the gap between STE and the other methods (see down panel in Figure 6).

B.2 "2 moons" example

We consider the binary classification problem "2 moons dataset" presented in Section 4 and inspired by Meng et al. (2020). The training dataset consists of $n = 2000$ samples and 200 test samples and is displayed in Figure 7a. A BNN with 9 weights is trained with one-hot coding and logistic loss. This BNN uses ReLu activations and its architecture is shown in Figure 7b. Four gradient-based approaches - a full precision NN, BinaryConnect, AdaSTE, and ASkewSGD - are compared to exhaustive search. In the latter, all 2^9 binary configurations on the training and test sets are compared. Figure 8 shows that different configurations lead to near-optimal performance. It is worth noting that permutation invariance implies that many solutions are equivalent in this simple example.

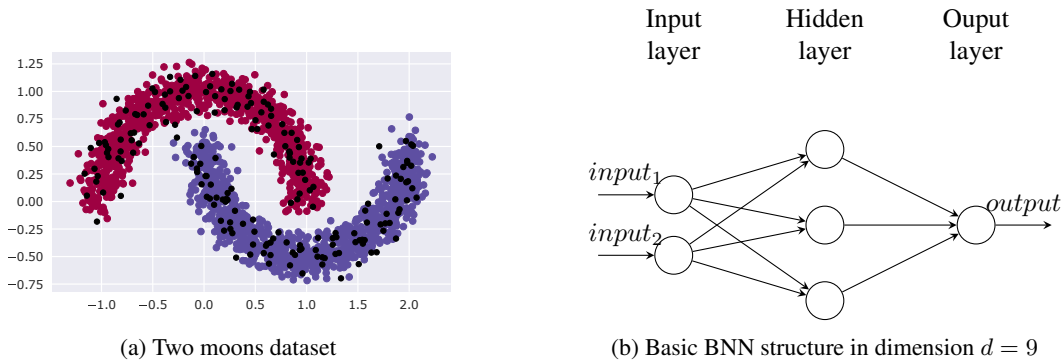


Figure 7: 2D Dataset and the associated BNN.




Figure 8: Histogram of test accuracies for the exhaustive search in dimension $d = 512$

B.3 Deep learning experiments

The performances reported in Section 4 were obtained with the best combination of hyperparameters that we tested. Other combinations are listed in Table 5 for the CIFAR-10 dataset. The performances reported in Table 2 are still dependent on

Table 5: Best Test accuracy after 100 training epochs on CIFAR-10.

α / lr	0.5/0.06	0.2/0.01	0.2/0.03	0.2/0.05	0.4/0.01
<code>ASkewSGD</code>	88.51	85.60	88.42	88.32	84.50

hyperparameter grid search and could be further improved if more resources are available.

B.4 ImageNet with binary weights

In this section, we compare the performance of `ASkewSGD` [W1/A32] on a large dataset and compare it to BinaryConnect Courbariaux et al. (2015); Hubara et al. (2016), Mirror Descent Ajanthan et al. (2021), AdaSTE Le et al. (2021), a standard full-precision NN, and a hypersphere-projected full-precision NN. To ensure a fair comparison, we compare the different methods using the same NN architecture. Moreover, we do not add bias in any layer, but introduce batch normalisation (without learning parameters) after each layer. The last connected layer is kept in full precision - a standard practice in BNN -. Contrary to (Liu et al., 2020; Chmiel et al., 2021), we have kept the first convolutional layer binary. We do not use layerwise scalar contrary to Rastegari et al. (2016).

We use a training setting similar to Section 4. We have adapted the code of Ajanthan et al. (2021); Le et al. (2021) to Resnet-18 for ImageNet. The hyperparameters for AdaSTE and MD are those prescribed for TinyImageNet. We use the same default data normalizations as the methods we compare to: we resize the input images to 256×256 and then randomly crop them to 224×224 while centering them to the appropriate sizes during training. Standard multiclass cross entropy loss is used. All models are fine-tuned for 100 epochs using the Adam (Kingma and Ba, 2014) optimizer with dynamics of 0.9 and 0.999 and a batch of size 512. The full precision NN is trained with an initial learning rate of 0.08. The projected full

precision NN uses a projected gradient algorithm. The same hyperparameters are used as in the "simple" algorithm, except that a deterministic projection onto the hypersphere is performed at each iteration. The ASkewSGD method is described in Algorithm 1, and we have set α to 0.5. The precision threshold ϵ is decreased from epoch to epoch: it is set to 1 at the beginning and then exponentially annealed to $.88^t$ in the last 50 epochs, where t is the epoch.

Just as with Section 4, we apply the function $\text{sign}(\cdot)$ to our NN before evaluating it on the test set. Each method was randomly initialized and independently executed once (due to ImageNet’s longer training time). The learning rate at epochs [20, 40] is divided by 2 for all methods. This task is more difficult than TinyImageNet’s, but we get the same result:

Table 6: Best Test accuracy after 100 training epochs.

Method	ImageNet (ResNet-18)	
	Top-1	Top-5
Full-precision	66.39	95.32
BinaryConnect	45.85	71.05
MD	46.38	71.18
AdaSTE	35.37	62.22
Projected gradient	2.58	7.93
ASkewSGD	46.95	72.11

ASkewSGD outperforms all current baselines. Moreover, ASkewSGD yields good results even when trained from scratch, compared to methods Bai et al. (2018); Liu et al. (2020) that require fine-tuning using a pre-trained network.

B.5 BNN with binary activations

BNN with binary weights and binary activations offer significant time savings in inference. We applied our training procedure ASkewSGD [W1/A1] to a VGG-small with $\text{sign}(\cdot)$ activations instead of ReLU activations to enable inference with only XNOR and bit-counting operations. The quantization of activations is here too extreme to apply the same procedure as in Section 4. The biased quantizer SAWB from Choi et al. (2018) does not work anymore (empirically). We assume the loss of neural gradient information is too important when activations are quantized on 2 levels.

During the training phase, a batch normalisation layer is inserted before each sign activation to scale the variance. In the backward pass, the derivative of $\text{sign}(\cdot)$ is approximated by the derivative of the function $\tanh(\cdot)$. During inference, we can get rid of the batch normalization (only the empirical mean is conserved and added to the bias term) and compute only binary operations.

We compare test accuracy with the CIFAR-10 dataset, which consists of 50000 training images and 10000 test images (in 10 classes). BNNs are fine-tuned for 100 epochs using the Adam optimizer with a dynamic range of 0.9 and 0.999 and a batch size of 100 with a learning rate of 0.03. The best test classification accuracies of binary networks obtained with ASkewSGD are listed in Table 7 for different values of $\alpha \in [0.2, 0.5, 0.7]$. The preliminary results reported in Table 7 show

Table 7: Best Test accuracy after 100 training epochs.

α	0.2	0.5	0.7
ASkewSGD [W1/A1]	81.12	84.34	82.92

that ASkewSGD is state-of-the art for training BNNs with binary weights and activations. Activation with the function $\text{sign}(\cdot)$ leads to a loss in expressive power and consequently a loss in performance. Several works introduce additional tricks such as real scaling factors Rastegari et al. (2016) to bridge the gap between binary signals and their real counterparts. These tricks can be easily implemented in our approach ASkewSGD.