
An Integer Programming Approach to Deep Neural Networks with Binary Activation Functions

Jannis Kurtz¹ Bubacarr Bah^{2,3}

Abstract

We study deep neural networks with binary activation functions (BDNN), i.e. the activation function only has two states. We show that the BDNN can be reformulated as a mixed-integer linear program which can be solved to global optimality by classical integer programming solvers. Additionally, a heuristic solution algorithm is presented and we study the model under data uncertainty, applying a two-stage robust optimization approach. We implemented our methods on random and real datasets and show that the heuristic version of the BDNN outperforms classical deep neural networks on the Breast Cancer Wisconsin dataset while performing worse on random data.

1. Introduction

Deep learning (DL) methods have of-late reinvigorated interest in artificial intelligence and data science, and they have had many successful applications in computer vision, natural language processing, and data analytics (LeCun et al., 2015). The training of deep neural networks relies mostly on (stochastic) gradient descent, hence the use of differentiable activation functions like ReLU, sigmoid or the hyperbolic tangent is the state of the art (Rumelhart et al., 1986; Goodfellow et al., 2016). On the contrary discrete activations, which may be more analogous to biological activations, present a training challenge due to non-differentiability and even discontinuity. If additionally the weights are considered to be binary, the use of discrete activation networks can reduce the computation and storage complexities, provides for better interpretation of solutions, and has the potential

^{*}Equal contribution ¹Chair for Mathematics of Information Processing, RWTH Aachen University, Aachen, Germany ²African Institute for Mathematical Sciences, Cape Town, South Africa ³Division of Applied Mathematics, Stellenbosch University, Stellenbosch, South Africa. Correspondence to: Jannis Kurtz <kurtz@mathc.rwth-aachen.de>.

to be more robust to adversarial perturbations than the continuous activation networks (Qin et al., 2020). Furthermore low-powered computations may benefit from discrete activations as a form of coarse quantizations (Plagianakos et al., 2001; Bengio et al., 2013; Courbariaux et al., 2015; Rastegari et al., 2016). Nevertheless, gradient descent-based training behaves like a black box, raising a lot of questions regarding the explainability and interpretability of internal representations (Hampson & Volper, 1990; Plagianakos et al., 2001; Bengio et al., 2013).

On the other hand, integer programming (IP) is known as a powerful tool to model a huge class of real-world optimization problems (Wolsey, 1998). Recently it was successfully applied to machine learning problems involving sparsity constraints and to evaluate trained neural networks (Bertsimas et al., 2017; 2019b; Fischetti & Jo, 2018).

Contributions: In Section 3, we reformulate the binary neural network as an IP formulation, referred to in the sequel as a binary deep neural network (BDNN), which can be solved to global optimality by classical IP solvers. All our results can also be applied to BDNNs with binary weights. Solving the latter class of networks to optimality was still an open problem. We note that it is straight forward to extend the IP formulation to more general settings and several variations of the model. We present a heuristic algorithm to calculate local optimal solutions of the BDNN in Section 3.1 and apply a two-stage robust optimization approach to protect the model against adversarial attacks in Section 4. In Section 5 we present computational results for random and real datasets and compare the BDNN to a deep neural network using ReLU activations (DNN). Despite scalability issues and a slightly worse accuracy on random datasets, the results indicate that the heuristic version outperforms the DNN on the *Breast Cancer Wisconsin* dataset.

2. Related Literature

The interest in BDNN goes back to (McCulloch & Pitts, 1943) where BDNNs were used to simulate Boolean functions. However, until the beginning of this century concerted efforts were made to train these networks either by specific schemes (Gray & Michel, 1992; Kohut & Stein-

bach, 2004) or via back propagation by modifications of the gradient descent method (Widrow & Winter, 1988; Toms, 1990; Barlett & Downs, 1992; Goodman & Zeng, 1994; Corwin et al., 1994; Plagianakos et al., 2001). More recent work regarding the back propagation, mostly motivated by the low complexity of computation and storage, build-on the pioneering works of (Bengio et al., 2013; Courbariaux et al., 2015; Hubara et al., 2016; Rastegari et al., 2016; Kim & Smaragdis, 2016); see (Qin et al., 2020) for a detailed survey. Regarding the generalization error of BDNN, it was already proved that the VC-Dimension of deep neural networks with binary activation functions is $\rho \log(\rho)$, where ρ is the number of weights of the BDNN; see (Baum & Haussler, 1989; Maass, 1994; Sakurai, 1993).

One the other hand IP methods were successfully applied to sparse classification problems (Bertsimas et al., 2017), sparse regression (Bertsimas et al., 2019b) and to evaluate trained neural networks (Fischetti & Jo, 2018). In (Icarte et al., 2019) BDNNs with weights restricted to $\{-1, 1\}$ are trained by a hybrid method based on constraint programming and mixed-integer programming. In (Khalil et al., 2018) the authors calculate optimal adversarial examples for BDNNs using a MIP formulation and integer propagation. Furthermore, robust optimization approaches were used to protect against adversarial attacks for other machine learning methods (Xu et al., 2009b;a; Bertsimas et al., 2019a).

3. Discrete Neural Networks

In this work we study *binary deep neural networks* (BDNN), i.e. classical deep neural networks with binary activation functions. As in the classical framework, for a given input vector $x \in \mathbb{R}^n$ we study classification functions f of the form $f(x) = \sigma^K (W^K \sigma^{K-1} (W^{K-1} \dots \sigma^1 (W^1 x) \dots))$ for weight matrices $W^k \in \mathbb{R}^{d_k \times d_{k-1}}$ and activation functions σ^k , which are applied component-wise. The dimension d_k is called the *width* of the k -th layer. In contrast to the recent developments of the field we consider the activation functions to be binary, more precisely each function is of the form

$$\sigma^k(\alpha) = \begin{cases} 0 & \text{if } \alpha < \lambda_k \\ 1 & \text{otherwise} \end{cases} \quad (1)$$

for $\alpha \in \mathbb{R}$ where the parameters $\lambda_k \in \mathbb{R}$ can be learned by our model simultaneously with the weight matrices which is normally not the case in the classical neural network approaches. Note that it is also possible to fix the values λ_k in advance.

In the following we use the notation $[p] := \{1, \dots, p\}$ for $p \in \mathbb{N}$.

Given a set of labeled training samples

$$X \times Y = \{(x^i, y^i) \mid i \in [m]\} \subset \mathbb{R}^n \times \{0, 1\}$$

we consider loss functions

$$\ell : \{0, 1\} \times \mathbb{R}^{d_K} \rightarrow \mathbb{R} \quad (2)$$

and the task is to find the optimal weight matrices which minimize the empirical loss over the training samples, i.e. we want to solve the problem

$$\begin{aligned} \min \quad & \sum_{i=1}^m \ell(y^i, z^i) \\ \text{s.t.} \quad & z^i = \sigma^K (W^K \sigma^{K-1} (\dots \sigma^1 (W^1 x^i) \dots)) \quad \forall i \in [m] \\ & W^k \in \mathbb{R}^{d_k \times d_{k-1}} \quad \forall k \in [K] \\ & \lambda_k \in \mathbb{R} \quad \forall k \in [K] \end{aligned} \quad (3)$$

for given dimensions d_0, \dots, d_K where $d_0 = n$. We set $d_K =: L$. All results in this work even hold for *regression problems*, more precisely for loss functions

$$\ell_r : \mathbb{R}^L \times \mathbb{R}^L \rightarrow \mathbb{R}$$

where we minimize the empirical loss $\sum_{i=1}^m \ell_r(x^i, z^i)$ instead. We use labels $\{0, 1\}$ instead of the classical labels $\{-1, 1\}$ due to ease of notation in our model. Nevertheless the same approach can be adapted to labels $\tilde{y}^i \in \{-1, 1\}$ by adding the additional constraints

$$\tilde{y}^i = -1 + 2y^i.$$

Several variants of the model can be considered:

- The output of the neural network is 0 or 1, defining the class the input is assigned to. As a loss function, we count the number of misclassified training samples. This variant can be modeled by choosing the last layer to have width $d_K = 1$ together with the loss function $\ell(y, z) = |y - z|$.
- The output of the neural network is a continuous vector $z \in \mathbb{R}^L$. This case can be modeled by choosing the activation function of the last layer σ^K as the identity and $d_K = L$. Any classical loss functions maybe consider, e.g. the squared loss $\ell(x, z) = \|y - z\|_2^2$. Note that a special case of this model is the *autoencoder* where in contrast to the continuous framework the input is encoded into a vector with 0-1 entries.
- All results can easily be generalized to multiclass classification; see Section 3.

In the following lemma we show how to reformulate Problem (3) as an integer program.

Lemma 1. *Assume the euclidean norm of each data point in X is bounded by r , then Problem (3) is equivalent to the mixed-integer non-linear program*

$$\min \sum_{i=1}^m \ell(y^i, u^{i,K}) \quad s.t. \quad (4)$$

$$W^1 x^i < M_1 u^{i,1} + \lambda_1 \quad (5)$$

$$W^1 x^i \geq M_1(u^{i,1} - 1) + \lambda_1 \quad (6)$$

$$W^k u^{i,k-1} < M_k u^{i,k} + \lambda_k \quad \forall k \in [K] \setminus \{1\} \quad (7)$$

$$W^k u^{i,k-1} \geq M_k(u^{i,k} - 1) + \lambda_k \quad \forall k \in [K] \setminus \{1\} \quad (8)$$

$$W^k \in [-1, 1]^{d_k \times d_{k-1}}, \quad \lambda_k \in [-1, 1] \quad \forall k \in [K] \quad (9)$$

$$u^{i,k} \in \{0, 1\}^{d_k} \quad \forall k \in [K], i \in [m], \quad (10)$$

where $M_1 := (nr + 1)$ and $M_k := (d_{k-1} + 1)$.

Proof. First we show that, due to the binary activation functions, we may assume $W^k \in [-1, 1]^{d_k \times d_{k-1}}$ and $\lambda_k \in [-1, 1]$ for all $k \in [K]$ in our model. To prove this assume we have any given solution W^1, \dots, W^K and corresponding $\lambda_1, \dots, \lambda_K$ of problem (3) with arbitrary values in \mathbb{R} . Consider any fixed layer $k \in [K]$. The k -th layer receives a vector $h^{k-1} \in \{0, 1\}^{d_{k-1}}$ from the previous layer, which is applied to W^k and afterwards the activation function σ^k is applied component-wise, i.e. the output of the k -th layer is a vector

$$h_j^k = \begin{cases} 0 & \text{if } (w_j^k)^\top h^{k-1} < \lambda_k \\ 1 & \text{otherwise} \end{cases}$$

where w_j^k is the j -th row of the matrix W^k . Set

$$\beta := \max\{|\lambda_k|, \max_{j=1, \dots, d_k} |w_{jl}^k|\}$$

and define $\tilde{W}^k := \beta W^k$ and $\tilde{\lambda}_k := \beta \lambda_k$. Then replacing W^k by \tilde{W}^k and λ_k by $\tilde{\lambda}_k$ in the k -th layer yields the same output vector h^k , since the inequality $(w_j^k)^\top h^{k-1} < \lambda_k$ holds if and only if the inequality $(\tilde{w}_j^k)^\top h^{k-1} < \tilde{\lambda}_k$ holds. Furthermore all entries of \tilde{W}^k and $\tilde{\lambda}_k$ have values in $[-1, 1]$.

Next we show that the constraints (5)–(8) correctly model the equation

$$\begin{aligned} z^i &:= u^{i,K} \\ &= \sigma^K(W^K \sigma^{K-1}(W^{K-1} \dots \sigma^1(W^1 x^i) \dots)) \end{aligned}$$

of Problem (3). The main idea is that the $u^{i,k}$ -variables model the output of the activation functions of data point i in layer k , i.e. they have value 0 if the activation value is 0 or value 1 otherwise. More precisely for any solution W^1, \dots, W^K and $\lambda_1, \dots, \lambda^K$ of the Problem in Lemma 1 the variable $u_j^{i,1}$ is equal to 1 if and only if $(w_j^1)^\top x^i \geq \lambda_1$

since otherwise Constraint (6) would be violated. Note that if $u_j^{i,1} = 1$, then Constraint (5) is always satisfied since all entries of W^1 are in $[-1, 1]$ and all entries of x^i are in $[-r, r]$ and therefore $|W^1 x^i| \leq nr < M_1$. Similarly we can show that $u_j^{i,1} = 0$ if and only if $(w_j^1)^\top x^i < \lambda_1$. Hence $u^{i,1}$ is the output of the first layer for data point x^i which is applied to W^2 in Constraints (7) and (8). By the same reasoning applied to Constraints (7) and (8) we can show that $u^{i,k}$ is equal to the output of the k -th layer for data point x^i for each $k \in [K] \setminus \{1\}$. Note that instead of the value nr we can use d_{k-1} here since the entries of $u^{i,k-1}$ can only have values 0 or 1 and the dimension of the rows of W^k is d_{k-1} . \square

The formulation in Lemma 1 is a non-linear mixed-integer programming (MINLP) formulation, since it contains products of variables, where each is a product of a continuous variable and an integer variable. By linearizing these products we can prove the following theorem.

Theorem 2. *Under the assumption of Lemma 1 Problem (3) is equivalent to the mixed-integer program*

$$(MIP) : \min \sum_{i=1}^m \ell(y^i, u^{i,K}) \quad s.t.$$

$$W^1 x^i < M_1 u^{i,1} + \lambda_1 \quad \forall i \in [m]$$

$$W^1 x^i \geq M_1(u^{i,1} - 1) + \lambda_1 \quad \forall i \in [m]$$

$$\sum_{l=1}^{d_{k-1}} s_l^{i,k} < M_k u^{i,k} + \lambda_k$$

$$\forall i \in [m], k \in [K] \setminus \{1\}$$

$$\sum_{l=1}^{d_{k-1}} s_l^{i,k} \geq M_k(u^{i,k} - 1) + \lambda_k$$

$$\forall i \in [m], k \in [K] \setminus \{1\}$$

$$s_{lj}^{i,k} \leq u_j^{i,k}, \quad s_{lj}^{i,k} \geq -u_j^{i,k}$$

$$\forall i \in [m], k \in [K] \setminus \{1\}, l \in [d_{k-1}], j \in [d_k]$$

$$s_{lj}^{i,k} \leq w_{lj}^k + (1 - u_j^{i,k})$$

$$\forall i \in [m], k \in [K] \setminus \{1\}, l \in [d_{k-1}], j \in [d_k]$$

$$s_{lj}^{i,k} \geq w_{lj}^k - (1 - u_j^{i,k})$$

$$\forall i \in [m], k \in [K] \setminus \{1\}, l \in [d_{k-1}], j \in [d_k]$$

$$W^k \in [-1, 1]^{d_k \times d_{k-1}} \quad \forall k \in [K]$$

$$\lambda_k \in [-1, 1] \quad \forall k \in [K]$$

$$u^{i,k} \in \{0, 1\}^{d_k} \quad \forall k \in [K], i \in [m]$$

$$s_l^{i,k} \in [-1, 1]^{d_k} \quad \forall i \in [m], k \in [K] \setminus \{1\}, l \in [d_{k-1}]$$

Proof. We replace each product of variables $w_{lj}^k u_j^{i,k-1}$ in the formulation of Lemma (1) by a new variable

$s_{lj}^{i,k} \in [-1, 1]$. To ensure that

$$w_{lj}^k u_j^{i,k-1} = s_{lj}^{i,k}$$

holds, we have to add the set of inequalities

$$\begin{aligned} s_{lj}^{i,k} &\leq u_j^{i,k} \\ s_{lj}^{i,k} &\geq -u_j^{i,k} \\ s_{lj}^{i,k} &\leq w_{lj}^k + (1 - u_j^{i,k}) \\ s_{lj}^{i,k} &\geq w_{lj}^k - (1 - u_j^{i,k}). \end{aligned}$$

Note that if $u_j^{i,k} = 0$, then the first two constraints ensure that $s_{lj}^{i,k} = 0$. Since $w_{lj}^k \in [-1, 1]$ this combination is also feasible for the last two constraints. If $u_j^{i,k} = 1$, then the last two constraints ensure, that $s_{lj}^{i,k} = w_{lj}^k$, while $s_{lj}^{i,k}$ and $u_j^{i,k}$ are still feasible for the first two constraints. \square

Remark 3. *The regression variant can be modeled by replacing the $u^{i,K}$ variables by continuous variables $z^i \in \mathbb{R}^{d_K}$ for each $i \in [m]$, replacing the objective function (4) by*

$$\sum_{i=1}^m \ell_r(x^i, z^i)$$

and replacing Constraints (7)-(8) with $k = K$ by the constraints

$$z^i = W^K u^{i,k-1}.$$

Remark 4. *Besides others the following classical loss functions can be considered:*

- *To model the empirical error we simply define the loss function by*

$$\sum_{i=1}^m \mathbb{1}_{\{z^i \neq y^i\}} = \sum_{i=1: y^i=0}^m z^i + \sum_{i=1: y^i=1}^m (1 - z^i)$$

In this case we obtain a mixed-integer linear problem (MILP) in Theorem 2.

- *Theorem 2 shows that the regression variant with quadratic loss $\ell_r(x, z) = \|x - z\|_2^2$ can be modelled as a mixed-integer quadratic problem (MIQP).*

Problem (MIP) can be solved by any off-the-shelf solvers like Gurobi for many classical loss functions (Gurobi Optimization, 2020). Unfortunately, the number of integer variables of the MIP formulation is $\mathcal{O}(DKm)$, where D is the maximum dimension of the layers, and therefore grows linear with the number of training samples and with the number of layers. For practical applications requiring large training sets solving the MIP formulation can be a hard or even impossible task. To tackle these difficulties we propose a fast heuristic in Section 3.1. Despite the latter drawback the MIP formulation has a lot of advantages and can give further insights into the analysis of deep neural networks:

- The BDNN with binary weights can be modeled by simply setting the weight variables to be in $\{0, 1\}$. The linearization results of Theorem (2) apply also in this case.

- Multiclass classification tasks can be modeled by using integer variables $u_j^{i,K} \in \mathbb{Z} \cap [a, b]$ in the last layer.

- More general discrete activation functions of the form

$$\sigma^k(\alpha) = v \text{ if } \underline{\lambda}_k^v \leq \alpha \leq \overline{\lambda}_k^v, v \in V \subset \mathbb{Z}$$

for a finite set V and pairwise non-intersecting intervals $[\underline{\lambda}_k^v, \overline{\lambda}_k^v]$ can be modeled by adding copies $u^{i,k,v}$ of the $u^{i,k}$ variables for each $v \in V$ and adding the constraints

$$W^k \left(\sum_{v \in V} v u^{i,k-1,v} \right) \leq M_k (1 - u^{i,k,v}) + \lambda_k^v$$

$$W^k \left(\sum_{v \in V} v u^{i,k-1,v} \right) \geq -M_k (1 - u^{i,k,v}) + \lambda_k^v$$

for each $v \in V$, $k \in [K] \setminus \{1\}$ and $i \in [m]$. The two constraints for the first layer are defined similarly, replacing $(\sum_{v \in V} v u^{i,k-1,v})$ by x^i . Note that the values $\lambda_k^v, \overline{\lambda}_k^v$ either have to be fixed in advance for each $v \in V$ or additional constraints have to be added which ensure the interval structure.

- The MIP formulation can easily be adjusted for applications where further constraints are desired. E.g. sparsity constraints of the form

$$\|W^k\|_0 \leq s$$

can be easily added to the formulation. Here $\|W^k\|_0$ is the number of non-zero entries of W^k .

- Any classical approaches handling uncertainty in the data can be applied to the MIP formulation. In Section 4 we will apply a robust optimization approach to the MIP formulation.

- The model is very flexible regarding changes in the training set. To add new data points that were not yet considered we just have to add the corresponding variables and constraints for the new data points to our already existing model and restart a solver, which is based on the idea of online machine learning.

- Classical solvers like Gurobi use branch & bound methods to solve MIP formulations. During these methods at each time the optimality gap, i.e. the percental difference between the best known upper and lower bound, is known. These methods can be stopped after a certain optimality gap is reached.

3.1. Heuristic Algorithm

In this section, we present a heuristic algorithm that is based on local search applied to the non-linear formulation in Lemma 1. This method is also known under the name *Mountain-Climbing method*; see (Nahapetyan, 2009). The idea is to avoid the quadratic terms by alternately optimizing the MINLP formulation over a subset of the variables and afterward over the complement of variables. Since for given weight variables exactly one feasible solution for the u -variables exists, this procedure would terminate after one iteration if all u -variables are contained in one of the problems. To avoid this problem we go through all layers and alternately fix the u or the W -variables. The second problem uses exactly the opposite variables for the fixation. More precisely we iteratively solve the following two problems:

$$\begin{aligned}
 & \min \sum_{i=1}^m \ell(y^i, u^{i,K}) \\
 \text{s.t. } & W^1 x^i < M_1 u^{i,1} + \lambda_1 \quad \forall i \in [m] \\
 & W^1 x^i \geq M_1 (u^{i,1} - 1) + \lambda_1 \quad \forall i \in [m] \\
 & W^k u^{i,k-1} < M_k u^{i,k} + \lambda_k \\
 & \quad \forall i \in [m], k \in [K] \setminus \{1\} \\
 & W^k u^{i,k-1} \geq M_k (u^{i,k} - 1) + \lambda_k \quad (H1) \\
 & \quad \forall i \in [m], k \in [K] \setminus \{1\} \\
 & W^k \in [-1, 1]^{d_k \times d_{k-1}}, \lambda_k \in [-1, 1] \\
 & \quad \forall k \in [K] : k \text{ odd} \\
 & u^{i,k} \in \{0, 1\}^{d_k} \\
 & \quad \forall i \in [m], k \in [K-1] : k \text{ odd} \\
 & u^{i,K} \in \{0, 1\}^{d_K} \quad \forall i \in [m].
 \end{aligned}$$

and

$$\begin{aligned}
 & \min \sum_{i=1}^m \ell(y^i, u^{i,K}) \\
 \text{s.t. } & W^1 x^i < M_1 u^{i,1} + \lambda_1 \quad \forall i \in [m] \\
 & W^1 x^i \geq M_1 (u^{i,1} - 1) + \lambda_1 \quad \forall i \in [m] \\
 & W^k u^{i,k-1} < M_k u^{i,k} + \lambda_k \\
 & \quad \forall i \in [m], k \in [K] \setminus \{1\} \\
 & W^k u^{i,k-1} \geq M_k (u^{i,k} - 1) + \lambda_k \quad (H2) \\
 & \quad \forall i \in [m], k \in [K] \setminus \{1\} \\
 & W^k \in [-1, 1]^{d_k \times d_{k-1}}, \lambda_k \in [-1, 1] \\
 & \quad \forall k \in [K] : k \text{ even} \\
 & u^{i,k} \in \{0, 1\}^{d_k} \\
 & \quad \forall i \in [m], k \in [K-1] : k \text{ even} \\
 & u^{i,K} \in \{0, 1\}^{d_K} \quad \forall i \in [m].
 \end{aligned}$$

In Problem (H1) all variables in

$$V_{\text{fix}}^1 := \{W^k, \lambda_k, u^{i,k} \text{ where } k \neq K \text{ and } k \text{ is even}\}$$

are fixed to given values; while in Problem (H2) all variables in

$$V_{\text{fix}}^2 := \{W^k, \lambda_k, u^{i,k} \text{ where } k \neq K \text{ and } k \text{ is odd}\}$$

are fixed to given values. In the heuristic procedure the fixed variables are always set to the optimal value of the preceding problem. Note that both problems are linear mixed-integer problems with roughly half of the variables of Problem (MIP). The heuristic is shown in Algorithm 1. Note that Algorithm 1 terminates after a finite number of steps, since there only exist finitely many possible variable assignments for the variables $u^{i,K}$, and therefore only finitely many objective values exist.

Algorithm 1 (Local Search Heuristic)

Input: $X \times Y, K, d_0, \dots, d_K$

Output: Weights W^1, \dots, W^K .

Draw random values for the variables in V_{fix}^1

repeat

 Calculate optimal solution of (H1), if feasible, for the current fixations in V_{fix}^1 .

 Set the values of all variables in V_{fix}^2 to the corresponding optimal values of (H1).

 Calculate optimal solution of (H2) if feasible for the current fixations in V_{fix}^2 .

 Set the values of all variables in V_{fix}^1 to the corresponding optimal values of (H2).

until no better solution is found

Return: $W = \{W^k\}_{k \in [K]}$

4. Data Uncertainty

In this section we consider labeled data

$$X \times Y = \{(x^i, y^i) \mid i \in [m]\}$$

where the euclidean norm of each data point is bounded by $r > 0$ and the data points are subject to uncertainty, i.e. a true data point $x^i \in \mathbb{R}^n$ can be disturbed by an unknown deviation vector $\delta \in \mathbb{R}^n$. One approach to tackle data uncertainty is based on the idea of robust optimization and was already studied for regression problems and support vector machines in (Bertsimas et al., 2019a; Xu et al., 2009b;a). In the robust optimization setting, we assume that for each data point x^i we have a convex set of possible deviation vectors $U^i \subset \mathbb{R}^n$ called *uncertainty set* which is defined by

$$U^i = \{\delta \in \mathbb{R}^n \mid \|\delta\| \leq r_i\}$$

for a given norm $\|\cdot\|$ and radii r_1, \dots, r_m . Note that classical convex sets like boxes, polyhedrons, or ellipsoids can

be modeled as above by using the ℓ_1 or ℓ_2 norm. The task is to find the weights of a neural network which are robust against all possible data disturbances in the uncertainty set $U := U^1 \times \dots \times U^m$, i.e. we want to find weights W^1, \dots, W^k which minimize the worst-case loss over the training set and which are feasible for all possible disturbances. The difference to the situations studied in (Bertsimas et al., 2019a; Xu et al., 2009b;a) is that the variables $u^{i,k}$ in our model, which are representing the value of the activation function, should be determined after the disturbance is known since the disturbance influences the data point and therefore the value of the activation function. Hence we can interpret the decision as a function of the uncertain parameters, i.e. $u^{i,k} : U^i \rightarrow \{0, 1\}$. A useful approach to tackle this situation is called two-stage robustness (or adaptive robustness) and was already studied intensively in the literature (Buchheim & Kurtz, 2018). Applied to our model the two-stage robust counterpart is

$$\min_{W^k, \lambda_k} \max_{\delta \in U} \min_{u^{i,k}} \left\{ \sum_{i=1}^m \ell(y^i, u^{i,K}) : W^1, \dots, W^k, \right. \\ \left. \lambda_1, \dots, \lambda_k, u^{1,1}, \dots, u^{m,K} \in P(X, \delta) \right\}. \quad (11)$$

where $P(X, \delta)$ is the set of feasible solutions of the inequality system

$$\begin{aligned} W^1(x^i + \delta^i) &< M_1^i u^{i,1} + \lambda_1 \quad \forall i \in [m] \\ W^1(x^i + \delta^i) &\geq M_1^i (u^{i,1} - 1) + \lambda_1 \quad \forall i \in [m] \\ W^k u^{i,k-1} &< M_k u^{i,k} + \lambda_k \quad \forall i \in [m], k \in [K] \setminus \{1\} \\ W^k u^{i,k-1} &\geq M_k (u^{i,k} - 1) + \lambda_k \quad \forall i \in [m], k \in [K] \setminus \{1\} \\ W^k &\in [-1, 1]^{d_k \times d_{k-1}}, \lambda_k \in [-1, 1] \quad \forall k \in [K] \\ u^{i,k} &\in \{0, 1\}^{d_k} \quad \forall i \in [m], k \in [K] \end{aligned}$$

and $M_1^i := n(r + r_i)$. The only difference to Constraints (5) – (10) is the appearance of the disturbance vectors δ^i in the first layer.

The idea is that in the first-stage the variables W^k and λ^k have to be calculated before the disturbance δ is known. For each possible first-stage solution the worst-case objective value overall scenarios $(\delta^1, \dots, \delta^m) \in U$ is considered in the objective function. In the second-stage, after the scenario is known, the best second-stage decision for the variables $u^{i,k}$ is made. Since for each combination of first-stage variables and scenarios in U exactly one second-stage variable is feasible, the second-stage minimization problem is equivalent to just finding this unique feasible solution, i.e. checking for each neuron if it is activated under the given disturbance or not. Problems of this form can be solved by so-called *column-and-constraint algorithms* (Zeng & Zhao,

2013). The idea of this method is to start with a finite subset of scenarios in U and iteratively add new scenarios to the problem. The new scenario is determined by an adversarial problem. This scenario is then added to the master-problem together with a copy of all second-stage variables, which will define the second-stage reaction to the new scenario. In this method, an increasing lower bound of Problem (11) is given by the master-problem, while the adversarial problem provides an upper bound to the problem. Therefore in each iteration, the method provides an upper bound on the optimality gap. Unfortunately, this method performs very bad for discrete optimization problems even if the uncertainty only affects the objective function (Kämmerling & Kurtz, 2020). Since to date no appropriate alternative methods exist, algorithmic progress in the field of two-stage robustness is desired to solve Problem (11) on realistic data sets.

5. Computations

In this section, we computationally compare the BDNN to the classical DNN with the ReLU activation function. We study networks with one hidden layer of dimension d_1 . All solution methods were implemented in Python 3.8 on an Intel(R) Core(TM) i5-4460 CPU with 3.20GHz and 8 GB RAM. The classical DNN was implemented by using the Keras API where we used the ReLU activation function on the hidden layer and the Softmax on the output layer. We used the binary cross entropy loss function. The number of epochs was set to 100. The exact IP formulation is given in Theorem 2 and all IP formulations used in the local search heuristic were implemented in Gurobi 9.0 with standard parameter settings. The strict inequalities in the IP formulations were replaced by non-strict inequalities adding -0.0001 to the right-hand-side. For the IP formulations, we set a time limit (wall time) of 24 hours.

We generated 10 random datasets in dimension $N = 100$ each with $m = 200$ data points. The entries of the data points were drawn from a uniform distribution with values in $[0, 10]$ for one-third of the data points, having label 1, and with values in $[-10, 0]$ for the second third of the data points, having label 0. The remaining data points were randomly drawn with entries in $[-1, 1]$ and have randomly assigned labels. We split each dataset into a training set of 160 samples and a testing set of 40 samples. All computations were implemented for neural networks with one hidden layer of dimension $d_1 \in \{20, 50, 100\}$. Figure 1 shows the average classification accuracy on the testing set over all 10 datasets achieved by the methods trained on m of the training points. The results indicate that the exact and the heuristic version of the BDNN have lower accuracy than the DNN. Furthermore, the performance of both BDNN methods seem to be much more unstable and depend more on the choice of the training set. Interestingly for a hidden layer of dimension




Figure 1. Average accuracy over 10 random instances for networks with one hidden layer of dimension d_1 trained on m data points.




Figure 2. Average runtime (logarithmic scale) over 10 random instances for networks with one hidden layer of dimension d_1 trained on m data points.

100, the heuristic method performs much better and can even compete with the classical DNN. In Figure 2 we show the runtime of all methods over m . Clearly, the runtime of the BDNN methods is much higher and seems to increase linearly in the number of data points. For real-world data sets with millions of data points, this method will fail using state-of-the-art solvers. Surprisingly, the runtime of the heuristic algorithm seems to be nearly the same as for the exact version, while the accuracy can be significantly better.

Additionally, we study all methods on the Breast Cancer Wisconsin dataset (BCW) (Dua & Graff, 2017). Here we also test the BDNN version where the values λ_k are all set to 0 instead of being part of the variables; we indicate this

version by BDNN_0 . The dataset has $m = 699$ data points with $N = 9$ attributes which we split into 80% training data and 20% testing data. Again all computations were implemented for neural networks with one hidden layer of dimension $d_1 \in \{25, 50\}$. In Tables 1 and 2 we show the accuracy, precision, recall and F1 score of all methods on the BCW dataset for a fixed shuffle of the data returned by the scikit-learn method *train_test_split* with the seed set to 42. It turns out that the exact BDNN performs better if the values λ_k are set to 0, while the heuristic version performs better if the λ_k are trained. The heuristic version of the BDNN has the best performance for $d_1 = 25$, significantly

better than the DNN. For $d_1 = 50$ the DNN is slightly better. Nevertheless, the best accuracy of 95% for the BCW dataset was achieved by the heuristic BDNN for $d_1 = 25$. In Table 3 we compare the heuristic BDNN to the DNN on 10 random shuffles of the BCW dataset and record the average, maximum, and minimum accuracies over all 10 shuffles. It turns out that the heuristic BDNN outperforms the DNN with the best accuracy of 97.1%.

No results are reported for the robust BDNN version, presented in Problem (11) since the column-and-constraint algorithm was not able to terminate during 24 hours on the random instances even for a network with $d_1 = 20$ and 20 data points. Further future research regarding more efficient optimization methods for discrete robust two-stage problems is necessary to solve Problem (11).

Table 1. Performance on the *Breast Cancer Wisconsin* dataset.

METHOD	d_1	ACC. (%)	OPT. GAP (%)
BDNN	25	69.3	0.0
BDNN ₀	25	83.6	0.0
BDNN HEUR.	25	95.0	0.0
BDNN ₀ HEUR.	25	30.0	0.0
DNN	25	91.4	-
BDNN	50	69.3	0.0
BDNN ₀	50	84.3	0.51
BDNN HEUR.	50	89.3	0.0
BDNN ₀ HEUR.	50	71.4	0.0
DNN	50	91.4	-

Table 2. Performance on the *Breast Cancer Wisconsin* dataset.

METHOD	d_1	PREC. (%)	REC. (%)	F1 (%)
BDNN	25	48.0	69.3	56.7
BDNN ₀	25	83.2	83.6	83.1
BDNN HEUR.	25	95.0	95.0	95.0
BDNN ₀ HEUR.	25	38.7	30.0	17.6
DNN	25	91.8	91.4	91.5
BDNN	50	63.6	69.3	58.0
BDNN ₀	50	86.4	84.3	84.7
BDNN HEUR.	50	91.5	89.3	89.6
BDNN ₀ HEUR.	50	74.6	71.4	72.3
DNN	50	91.4	91.4	91.3

Table 3. Accuracy on the *Breast Cancer Wisconsin* dataset over 10 random shuffles of the data.

METHOD	d_1	AVG. (%)	MAX (%)	MIN (%)
BDNN HEUR.	25	93.2	97.1	85.0
DNN	25	89.1	91.4	85.7

6. Conclusion

We show that binary deep neural networks can be modeled by mixed-integer programming formulations, which can be solved to global optimality by classical integer programming solvers. Additionally, we present a heuristic algorithm and a robust version of the model. Our tests on random and real datasets indicate that binary deep neural networks are much more unstable and computationally harder to solve than classical DNNs. Nevertheless, the heuristic algorithm outperforms the classical DNN on the Breast Cancer Wisconsin dataset. Moreover, the mixed-integer programming formulation is very adjustable to variations of the model and could give new insights into the understanding of deep neural networks. This motivates further research into the scalability of the IP method and the study of other tractable reformulations of DL algorithms.

References

- Barlett, P. L. and Downs, T. Using random weights to train multilayer networks of hard-limiting units. *IEEE Transactions on Neural Networks*, 3(2):202–210, 1992.
- Baum, E. B. and Haussler, D. What size net gives valid generalization? In *Advances in neural information processing systems*, pp. 81–90, 1989.
- Bengio, Y., Léonard, N., and Courville, A. Estimating or propagating gradients through stochastic neurons for conditional computation. *arXiv preprint arXiv:1308.3432*, 2013.
- Bertsimas, D., Pauphilet, J., and Van Parys, B. Sparse classification: a scalable discrete optimization perspective. *arXiv preprint arXiv:1710.01352*, 2017.
- Bertsimas, D., Dunn, J., Pawlowski, C., and Zhuo, Y. D. Robust classification. *INFORMS Journal on Optimization*, 1(1):2–34, 2019a.
- Bertsimas, D., Pauphilet, J., and Van Parys, B. Sparse regression: Scalable algorithms and empirical performance. *arXiv preprint arXiv:1902.06547*, 2019b.
- Buchheim, C. and Kurtz, J. Robust combinatorial optimization under convex and discrete cost uncertainty. *EURO Journal on Computational Optimization*, 6(3):211–238, 2018.
- Corwin, E. M., Logar, A. M., and Oldham, W. J. An iterative method for training multilayer networks with threshold functions. *IEEE Transactions on Neural Networks*, 5(3):507–508, 1994.
- Courbariaux, M., Bengio, Y., and David, J.-P. Binaryconnect: Training deep neural networks with binary weights during propagations. In *Advances in neural information processing systems*, pp. 3123–3131, 2015.
- Dua, D. and Graff, C. UCI machine learning repository, 2017. URL <http://archive.ics.uci.edu/ml>.
- Fischetti, M. and Jo, J. Deep neural networks and mixed integer linear optimization. *Constraints*, 23(3):296–309, 2018.

- Goodfellow, I., Bengio, Y., and Courville, A. *Deep learning*. MIT press, 2016.
- Goodman, R. M. and Zeng, Z. A learning algorithm for multi-layer perceptrons with hard-limiting threshold units. In *Proceedings of IEEE Workshop on Neural Networks for Signal Processing*, pp. 219–228. IEEE, 1994.
- Gray, D. L. and Michel, A. N. A training algorithm for binary feedforward neural networks. *IEEE Transactions on Neural Networks*, 3(2):176–194, 1992.
- Gurobi Optimization, L. Gurobi optimizer reference manual, 2020. URL <http://www.gurobi.com>.
- Hampson, S. E. and Volper, D. J. Representing and learning boolean functions of multivalued features. *IEEE transactions on systems, man, and cybernetics*, 20(1):67–80, 1990.
- Hubara, I., Courbariaux, M., Soudry, D., El-Yaniv, R., and Bengio, Y. Binarized neural networks: Training neural networks with weights and activations constrained to +1 or-1. *arXiv preprint arXiv:1602.02830*, 2016.
- Icarte, R. T., Illanes, L., Castro, M. P., Cire, A. A., McIlraith, S. A., and Beck, J. C. Training binarized neural networks using MIP and CP. In *International Conference on Principles and Practice of Constraint Programming*, pp. 401–417. Springer, 2019.
- Kämmerling, N. and Kurtz, J. Oracle-based algorithms for binary two-stage robust optimization. *Computational Optimization and Applications*, pp. 1–31, 2020.
- Khalil, E. B., Gupta, A., and Dilkina, B. Combinatorial attacks on binarized neural networks. *arXiv preprint arXiv:1810.03538*, 2018.
- Kim, M. and Smaragdis, P. Bitwise neural networks. *arXiv preprint arXiv:1601.06071*, 2016.
- Kohut, R. and Steinbach, B. Boolean neural networks. *Transactions on Systems*, 2:420–425, 2004.
- LeCun, Y., Bengio, Y., and Hinton, G. Deep learning. *nature*, 521(7553):436–444, 2015.
- Maass, W. Perspectives of current research about the complexity of learning on neural nets. In *Theoretical advances in neural computation and learning*, pp. 295–336. Springer, 1994.
- McCulloch, W. S. and Pitts, W. A logical calculus of the ideas immanent in nervous activity. *The bulletin of mathematical biophysics*, 5(4):115–133, 1943.
- Nahapetyan, A. G. *Bilinear Programming*, pp. 279–282. Springer US, Boston, MA, 2009. ISBN 978-0-387-74759-0. doi: 10.1007/978-0-387-74759-0_48. URL https://doi.org/10.1007/978-0-387-74759-0_48.
- Plagianakos, V., Magoulas, G., Nouis, N., and Vrahatis, M. Training multilayer networks with discrete activation functions. In *IJCNN'01. International Joint Conference on Neural Networks. Proceedings (Cat. No. 01CH37222)*, volume 4, pp. 2805–2810. IEEE, 2001.
- Qin, H., Gong, R., Liu, X., Bai, X., Song, J., and Sebe, N. Binary neural networks: A survey. *Pattern Recognition*, pp. 107281, 2020.
- Rastegari, M., Ordonez, V., Redmon, J., and Farhadi, A. XNOR-Net: Imagenet classification using binary convolutional neural networks. In *European conference on computer vision*, pp. 525–542. Springer, 2016.
- Rumelhart, D. E., Hinton, G. E., and Williams, R. J. Learning representations by back-propagating errors. *nature*, 323(6088):533–536, 1986.
- Sakurai, A. Tighter bounds of the VC-dimension of three layer networks. In *Proceedings of the World Congress on Neural Networks*, volume 3, pp. 540–543. Erlbaum, 1993.
- Toms, D. Training binary node feedforward neural networks by back propagation of error. *Electronics letters*, 26(21):1745–1746, 1990.
- Widrow, B. and Winter, R. Neural nets for adaptive filtering and adaptive pattern recognition. *Computer*, 21(3):25–39, 1988.
- Wolsey, L. A. *Integer programming*, volume 52. John Wiley & Sons, 1998.
- Xu, H., Caramanis, C., and Mannor, S. Robust regression and LASSO. In *Advances in Neural Information Processing Systems*, pp. 1801–1808, 2009a.
- Xu, H., Caramanis, C., and Mannor, S. Robustness and regularization of support vector machines. *Journal of machine learning research*, 10(Jul):1485–1510, 2009b.
- Zeng, B. and Zhao, L. Solving two-stage robust optimization problems using a column-and-constraint generation method. *Operations Research Letters*, 41(5):457–461, 2013.