Dichotomize and Generalize: PAC-Bayesian Binary Activated Deep Neural Networks

Gaël Letarte
Université Laval
Canada
gael.letarte.1@ulaval.ca

Pascal Germain
Inria
France
pascal.germain@inria.fr

Benjamin Guedj
Inria and University College London
France and United Kingdom
benjamin.guedj@inria.fr

François Laviolette
Université Laval
Canada
francois.laviolette@ift.ulaval.ca

Abstract

We present a comprehensive study of multilayer neural networks with binary activation, relying on the PAC-Bayesian theory. Our contributions are twofold: (i) we develop an end-to-end framework to train a binary activated deep neural network, overcoming the fact that binary activation function is non-differentiable; (ii) we provide nonvacuous PAC-Bayesian generalization bounds for binary activated deep neural networks. Noteworthy, our results are obtained by minimizing the expected loss of an architecture-dependent aggregation of binary activated deep neural networks. The performance of our approach is assessed on a thorough numerical experiment protocol on real-life datasets.

1 Introduction

The remarkable practical successes of deep learning make the need for better theoretical understanding all the more pressing. The PAC-Bayesian theory has recently emerged as a fruitful framework to analyze generalization abilities of deep neural network. Inspired by precursor work of Langford and Caruana [2001], nonvacuous risk bounds for multilayer architectures have been obtained by Dziugaite and Roy [2017], Zhou et al. [2019]. Although informative, these results upper bound the prediction error of a (stochastic) neural network with perturbed weights, which is not the one used to predict in practice. Moreover, the analysis does not explicitly take into account the network architecture (number of layers, neurons per layer, type of activation function). This is to the notable exception of the work of Neyshabur et al. [2018] which provides a PAC-Bayesian analysis relying on the network architecture and the choice of ReLU activation function. The latter bound gives valuable insights on the generalization mechanism of neural networks (namely in terms of the spectral norms of the learned weight matrices), but their validity hold for some margin assumptions, and they are likely to be numerically vacuous.

We focus our study on deep neural networks with a sign activation function. We call such networks *binary activated multilayer* (BAM) networks. This specialization leads to nonvacuous generalization bounds which hold under the sole assumption that training samples are *iid*. We provide a PAC-Bayesian bound holding on the generalization error of a continuous aggregation of BAM networks. This leads to an original approach to train BAM networks, named PBGNet. The building block of PBGNet is brought by the specialization of PAC-Bayesian bounds to linear classifiers [Germain et al., 2009], that we adapt to deep neural networks. The term *binary neural networks* has been coined by Bengio [2009], and further studied in Hubara et al. [2016, 2017], Soudry et al. [2014]: it refers to

neural networks for which both the activation functions and the weights are binarized (in contrast with BAM networks). These architectures are motivated by the desire to reduce the computation and memory footprints of neural networks.

Our theory-driven approach is validated on real life datasets, showing competitive accuracy with tanh-activated multilayer networks, and providing nonvacuous generalization bounds.

Organisation of the paper. We formalize our framework and notation in Section 2, along with a presentation of the PAC-Bayes framework and its specialization to linear classifiers. Section 3 illustrates the key ideas we develop in the present paper, on the simple case of a two-layers neural network. This is then generalized to deep neural networks in Section 4. We present our main theoretical result in Section 5: a PAC-Bayesian generalization bound for binary activated deep neural networks, and the associated learning algorithm. Section 6 presents the numerical experiment protocol and results. The paper closes with avenues for future work in Section 7.

2 Framework and notation

We stand in the supervised binary classification setting: given a real input vector[1] $\mathbf{x} \in \mathbb{R}^{d_0}$, one wants to predict a label $y \in \{-1, 1\}$. Let us consider a neural network of L *fully connected* layers with a (binary) sign activation function: $\mathrm{sgn}(a) = 1$ if $a > 0$ and $\mathrm{sgn}(a) = -1$ otherwise.[2] We let d_k denote the number of neurons of the k^{th} layer, for $k \in \{1, \ldots, L\}$; d_0 is the input data point dimension, and $D := \sum_{k=1}^{L} d_{k-1} d_k$ is the total number of parameters. The output of the (deterministic) BAM network on an input data point $\mathbf{x} \in \mathbb{R}^{d_0}$ is given by

$$f_\theta(\mathbf{x}) = \mathrm{sgn}\big(\mathbf{W}_L \mathrm{sgn}\big(\mathbf{W}_{L-1} \mathrm{sgn}\big(\ldots \mathrm{sgn}(\mathbf{W}_1 \mathbf{x})\big)\big)\big), \tag{1}$$

where $\mathbf{W}_k \in \mathbb{R}^{d_k \times d_{k-1}}$ denotes the weight matrices. The network is thus parametrized by $\theta = \mathrm{vec}\big(\{\mathbf{W}_k\}_{k=1}^{L}\big) \in \mathbb{R}^D$. The i^{th} line of matrix $\mathbf{W}_k$ will be denoted $\mathbf{w}_k^i$. For binary classification, the BAM network final layer $\mathbf{W}_L \in \mathbb{R}^{1 \times d_{L-1}}$ has one line ($d_L = 1$), that is a vector $\mathbf{w}_L \in \mathbb{R}^{d_{L-1}}$, and $f_\theta : \mathbb{R}^{d_0} \to \{-1, 1\}$. We study the classification accuracy under the linear loss $\ell(y', y) := \frac{1}{2}(1 - yy')$.

2.1 Elements from the PAC-Bayesian theory

The Probably approximately correct (PAC) framework [introduced by Valiant, 1984] holds under the frequentist assumption that data is sampled in an *iid* fashion from a data distribution $\mathcal{D}$ over the input-output space. The learning algorithm observes a finite training sample $S = \{(\mathbf{x}_i, y_i)\}_{i=1}^{n} \sim \mathcal{D}^{\otimes n}$ and outputs a predictor $f : \mathbb{R}^{d_0} \to \{-1, 1\}$. We define $\mathcal{L}_\mathcal{D}(f)$ as the generalization loss on the data generating distribution $\mathcal{D}$, and $\widehat{\mathcal{L}}_S(f)$ as the empirical error on the training set, given by

$$\mathcal{L}_\mathcal{D}(f) = \mathop{\mathbf{E}}_{(\mathbf{x},y) \sim \mathcal{D}} \ell(f(\mathbf{x}), y), \quad \text{and} \quad \widehat{\mathcal{L}}_S(f) = \frac{1}{n} \sum_{i=1}^{n} \ell(f(\mathbf{x}_i), y_i).$$

PAC-Bayes allows considering the expected loss of an aggregation of predictors: considering a distribution Q (called the *posterior*) over a family of binary classifiers $\mathcal{F}$, one obtains PAC upper bounds on $\mathbf{E}_{f \sim Q} \mathcal{L}_\mathcal{D}(f)$. Under the linear loss, this is equivalent to the loss of the predictor $F_Q(\mathbf{x}) := \mathbf{E}_{f \sim Q} f(\mathbf{x})$, performing a Q-aggregation of all classifiers in $\mathcal{F}$. In other words, we may upper bound with an arbitrarily high probability the generalization loss $\mathcal{L}_\mathcal{D}(F_Q)$, by its empirical counterpart $\widehat{\mathcal{L}}_S(F_Q)$ and a complexity term, the Kullback-Leibler divergence between Q and a reference measure (called the *prior* distribution) P chosen independently of the training set S, given by $\mathrm{KL}(Q\|P) := \int \ln \frac{Q(\theta)}{P(\theta)} Q(\mathrm{d}\theta)$. Since the seminal works of Shawe-Taylor and Williamson [1997], McAllester [1999, 2003] and Catoni [2003, 2004, 2007], the celebrated PAC-Bayesian theorem has been declined in many forms [see Guedj, 2019, for a survey]. The following Theorems 1 and 2 will be useful in the sequel.

Theorem 1 (Seeger [2002]). *Given a prior P on $\mathcal{F}$, with probability at least $1 - \delta$ over $S \sim \mathcal{D}^{\otimes n}$,*

$$\text{for all } Q \text{ on } \mathcal{F}: \quad \mathrm{kl}\Big(\widehat{\mathcal{L}}_S(F_Q) \big\| \mathcal{L}_\mathcal{D}(F_Q)\Big) \leq \frac{\mathrm{KL}(Q\|P) + \ln \frac{2\sqrt{n}}{\delta}}{n}, \tag{2}$$

---

[1] Bold uppercase letters denote matrices, bold lowercase letters denote vectors.

[2] We consider the activation function as an *element-wise* operator when applied to vectors or matrices.

where $\mathrm{kl}(q\|p) := q \ln \frac{q}{p} + (1-q)\ln\frac{1-q}{1-p}$ is the Kullblack-Leibler divergence between Bernoulli distributions with probability of success p and q, respectively.

Theorem 2 (Catoni [2007]). *Given P on $\mathcal{F}$ and $C > 0$, with probability at least $1 - \delta$ over $S \sim \mathcal{D}^{\otimes n}$,*

$$\text{for all } Q \text{ on } \mathcal{F}: \quad \mathcal{L}_\mathcal{D}(F_Q) \leq \frac{1}{1 - e^{-C}} \left(1 - \exp\left(-C\,\widehat{\mathcal{L}}_S(F_Q) - \frac{\mathrm{KL}(Q\|P) + \ln\frac{1}{\delta}}{n} \right) \right). \quad (3)$$

From Theorems 1 and 2, we obtain PAC-Bayesian bounds on the *linear loss* of the Q-aggregated predictor F_Q. The obtained bounds can be turned into bounds on the *zero-one loss* with an extra 2 multiplicative factor (using the elementary inequality $\mathbb{1}[y' \neq y] \leq 2\ell(y', y)$).

2.2 Elementary building block: PAC-Bayesian learning of linear classifiers

The PAC-Bayesian specialization to linear classifiers has been proposed by Langford and Shawe-Taylor [2002], and used for providing tight generalization bounds and a model selection criteria [further studied by Ambroladze et al., 2006, Langford, 2005, Parrado-Hernández et al., 2012]. This paved the way to the PAC-Bayesian bound minimization algorithm of Germain et al. [2009], that learns a linear classifier $f_\mathbf{w}(\mathbf{x}) := \mathrm{sgn}(\mathbf{w} \cdot \mathbf{x})$, with $\mathbf{w} \in \mathbb{R}^d$. The strategy is to consider a Gaussian posterior $Q_\mathbf{w} := \mathcal{N}(\mathbf{w}, I_d)$ and a Gaussian prior $P_{\mathbf{w}_0} := \mathcal{N}(\mathbf{w}_0, I_d)$ over the space of all linear predictors $\mathcal{F}_d := \{f_\mathbf{v} | \mathbf{v} \in \mathbb{R}^d\}$ (where I_d denotes the $d \times d$ identity matrix). The posterior is used to define a linear predictor $f_\mathbf{w}$ and the prior may have been learned on previously seen data; a common uninformative prior being the null vector $\mathbf{w}_0 = \mathbf{0}$. With such parametrization, $\mathrm{KL}(Q_\mathbf{w}\|P_{\mathbf{w}_0}) = \frac{1}{2}\|\mathbf{w} - \mathbf{w}_0\|^2$. Moreover, the $Q_\mathbf{w}$-aggregated output can be written in terms of the Gauss error function $\mathrm{erf}(\cdot)$. In Germain et al. [2009], the erf function is introduced as a loss function to be optimized. Here we interpret it as the predictor output, to be in phase with our neural network approach. Likewise, we study the linear loss of an aggregated predictor instead of the *Gibbs risk* of a stochastic classifier. We obtain (explicit calculations are provided in Appendix A.1 for completeness)

$$F_\mathbf{w}(\mathbf{x}) := \underset{\mathbf{v} \sim Q_\mathbf{w}}{\mathbf{E}} f_\mathbf{v}(\mathbf{x}) = \mathrm{erf}\left(\frac{\mathbf{w} \cdot \mathbf{x}}{\sqrt{d}\|\mathbf{x}\|}\right), \quad \text{with } \mathrm{erf}(x) := \frac{2}{\sqrt{\pi}} \int_0^x e^{-t^2} dt. \quad (4)$$

Given a training set $S \sim \mathcal{D}^{\otimes n}$, Germain et al. [2009] propose to minimize a PAC-Bayes upper bound on $\mathcal{L}_\mathcal{D}(F_\mathbf{w})$ by gradient descent on the coordinates of $\mathbf{w}$. This approach is appealing as the bounds are valid uniformly for all $Q_\mathbf{w}$ (see Equations 2 and 3). In other words, the algorithm provides both a learned predictor and a generalization guarantee that is rigorously valid (under the *iid* assumption) even when the optimization procedure did not find the global minimum of the cost function (either because it converges to a local minimum, or early stopping is used). Germain et al. [2009] investigate the optimization of several versions of Theorems 1 and 2. The minimization of Theorem 1 generally leads to tighter bound values, but empirical studies show lowest accuracy as the procedure conservatively prevents overfitting. The best empirical results are obtained by minimizing Theorem 2 for a fixed hyperparameter C, selected by cross-validation. Minimizing Equation (3) amounts to minimizing

$$C\,n\,\widehat{\mathcal{L}}_S(F_\mathbf{w}) + \mathrm{KL}(Q_\mathbf{w}\|P_{\mathbf{w}_0}) = C\frac{1}{2}\sum_{i=1}^{n} \mathrm{erf}\left(-y_i \frac{\mathbf{w} \cdot \mathbf{x}_i}{\sqrt{d}\|\mathbf{x}_i\|}\right) + \frac{1}{2}\|\mathbf{w} - \mathbf{w}_0\|^2. \quad (5)$$

In their discussion, Germain et al. [2009] observe that the objective in Equation (5) is similar to the one optimized by the soft-margin Support Vector Machines [Cortes and Vapnik, 1995], by roughly interpreting the *hinge loss* $\max(0, 1 - yy')$ as a convex surrogate of the *probit loss* $\mathrm{erf}(-yy')$. Likewise, Langford and Shawe-Taylor [2002] present this parameterization of the PAC-Bayes theorem as a margin bound. In the following, we develop an original approach to neural networks based on a slightly different observation: the predictor output given by Equation (4) is reminiscent of the $\tanh$ activation used in classical neural networks. Therefore, as the linear *perceptron* is viewed as the *building block* of modern multilayer neural networks, the PAC-Bayesian specialization to binary classifiers is the cornerstone of our theoretical and algorithmic framework for BAM networks.

3 The simple case of a one hidden layer network

Let us first consider a network with one hidden layer of size d_1. Hence, this network is parameterized by weights $\theta = \mathrm{vec}(\{\mathbf{W}_1, \mathbf{w}_2\})$, with $\mathbf{W}_1 \in \mathbb{R}^{d_1 \times d_0}$ and $\mathbf{w}_2 \in \mathbb{R}^{d_1}$. Given an input $\mathbf{x} \in \mathbb{R}^{d_0}$, the

output of the network is

$$f_\theta(\mathbf{x}) = \text{sgn}\big(\mathbf{w}_2 \cdot \text{sgn}(\mathbf{W}_1\mathbf{x})\big). \tag{6}$$

Following Section 2, we consider an isotropic Gaussian posterior distribution centered in θ, denoted $Q_\theta = \mathcal{N}(\theta, I_D)$, over the family of all networks $\mathcal{F}_D = \{f_{\tilde\theta} \mid \tilde\theta \in \mathbb{R}^D\}$. Thus, the prediction of the Q_θ-aggregate predictor is given by $F_\theta(\mathbf{x}) = \mathbf{E}_{\tilde\theta \sim Q_\theta} f_{\tilde\theta}(\mathbf{x})$. Note that Dziugaite and Roy [2017], Langford and Caruana [2001] also consider Gaussian distributions over neural networks parameters. However, as their analysis is not specific to a particular activation function—experiments are performed with *typical* activation functions (sigmoid, ReLU)—the prediction relies on sampling the parameters according to the posterior. An originality of our approach is that, by studying the sign activation function, we can calculate the exact solution of $F_\theta(\mathbf{x})$, as detailed below.

3.1 Deterministic network

Prediction. To compute the value of $F_\theta(\mathbf{x})$, we first need to decompose the probability of each $\tilde\theta = \text{vec}(\{\mathbf{V}_1, \mathbf{v}_2\}) \sim Q_\theta$ as $Q_\theta(\tilde\theta) = Q_1(\mathbf{V}_1)Q_2(\mathbf{v}_2)$, with $Q_1 = \mathcal{N}(\mathbf{W}_1, I_{d_0 d_1})$ and $Q_2 = \mathcal{N}(\mathbf{w}_2, I_{d_0})$.

$$F_\theta(\mathbf{x}) = \int_{\mathbb{R}^{d_1 \times d_0}} Q_1(\mathbf{V}_1) \int_{\mathbb{R}^{d_1}} Q_2(\mathbf{v}_2)\text{sgn}(\mathbf{v}_2 \cdot \text{sgn}(\mathbf{V}_1\mathbf{x}))d\mathbf{v}_2 d\mathbf{V}_1$$

$$= \int_{\mathbb{R}^{d_1 \times d_0}} Q_1(\mathbf{V}_1) \, \text{erf}\left(\frac{\mathbf{w}_2 \cdot \text{sgn}(\mathbf{V}_1\mathbf{x})}{\sqrt{2}\|\text{sgn}(\mathbf{V}_1\mathbf{x})\|}\right) d\mathbf{V}_1 \tag{7}$$

$$= \sum_{\mathbf{s} \in \{-1,1\}^{d_1}} \text{erf}\left(\frac{\mathbf{w}_2 \cdot \mathbf{s}}{\sqrt{2d_1}}\right) \int \mathbb{1}[\mathbf{s} = \text{sgn}(\mathbf{V}_1\mathbf{x})] Q_1(\mathbf{V}_1) \, d\mathbf{V}_1 \tag{8}$$

$$= \sum_{\mathbf{s} \in \{-1,1\}^{d_1}} \text{erf}\left(\frac{\mathbf{w}_2 \cdot \mathbf{s}}{\sqrt{2d_1}}\right) \Psi_{\mathbf{s}}(\mathbf{x}, \mathbf{W}_1), \tag{9}$$

where, from $Q_1(\mathbf{V}_1) = \prod_{i=1}^{d_1} Q_1^i(\mathbf{v}_1^i)$ with $Q_1^i := \mathcal{N}(\mathbf{w}_1^i, I_{d_0})$, we obtain

$$\Psi_{\mathbf{s}}(\mathbf{x}, \mathbf{W}_1) := \prod_{i=1}^{d_1} \int_{\mathbb{R}^{d_0}} \mathbb{1}[s_i \, \mathbf{x} \cdot \mathbf{v}_1^i > 0] Q_1^i(\mathbf{v}_1^i) \, d\mathbf{v}_1^i = \prod_{i=1}^{d_1} \underbrace{\left[\frac{1}{2} + \frac{s_i}{2}\text{erf}\left(\frac{\mathbf{w}_1^i \cdot \mathbf{x}}{\sqrt{2}\|\mathbf{x}\|}\right)\right]}_{\psi_{s_i}(\mathbf{x}, \mathbf{w}_1^i)}. \tag{10}$$

Line (7) states that the output neuron is a linear predictor over the hidden layer's activation values $\mathbf{s} = \text{sgn}(\mathbf{V}_1\mathbf{x})$; based on Equation (4), the integral on $\mathbf{v}_2$ becomes $\text{erf}\big(\mathbf{w}_2 \cdot \mathbf{s}/(\sqrt{2}\|\mathbf{s}\|)\big)$. As a function of $\mathbf{s}$, the latter expression is piecewise constant. Thus, line (8) discretizes the integral on $\mathbf{V}_1$ as a sum of the 2^{d_1} different values of $\mathbf{s} = (s_i)_{i=1}^{d_1}, s_i \in \{-1, 1\}$. Note that $\|\mathbf{s}\|^2 = d_1$.

Finally, one can compute the exact output of $F_\theta(\mathbf{x})$, provided one accepts to compute a sum combinatorial in the number of hidden neurons (Equation 9). We show in forthcoming Section 3.2 that it is possible to circumvent this computational burden and approximate $F_\theta(\mathbf{x})$ by a sampling procedure.

Derivatives. Following contemporary approaches in deep neural networks [Goodfellow et al., 2016], we minimize the empirical loss $\widehat{\mathcal{L}}_S(F_\theta)$ by stochastic gradient descent (SGD). This requires to compute the partial derivative of the cost function according to the parameters θ:

$$\frac{\partial \widehat{\mathcal{L}}_S(F_\theta)}{\partial \theta} = \frac{1}{n}\sum_{i=1}^n \frac{\partial \ell(F_\theta(\mathbf{x}_i), y_i)}{\partial \theta} = \frac{1}{n}\sum_{i=1}^n \frac{\partial F_\theta(\mathbf{x}_i)}{\partial \theta} \, \ell'(F_\theta(\mathbf{x}_i), y_i), \tag{11}$$

with the derivative of the linear loss $\ell'(F_\theta(\mathbf{x}_i), y_i) = -\frac{1}{2}y$. The partial derivatives of the prediction function (Equation 9) according to the hidden layer parameters $\mathbf{w}_1^k \in \{\mathbf{w}_1^1, \ldots, \mathbf{w}_1^{d_1}\}$ and the output neuron parameters $\mathbf{w}_2$ are

$$\frac{\partial}{\partial \mathbf{w}_1^k} F_\theta(\mathbf{x}) = \frac{\mathbf{x}}{2^{\frac{3}{2}}\|\mathbf{x}\|}\text{erf}'\left(\frac{\mathbf{w}_1^k \cdot \mathbf{x}}{\sqrt{2}\|\mathbf{x}\|}\right) \sum_{\mathbf{s} \in \{-1,1\}^{d_1}} s_k \, \text{erf}\left(\frac{\mathbf{w}_2 \cdot \mathbf{s}}{\sqrt{2d_1}}\right) \left[\frac{\Psi_{\mathbf{s}}(\mathbf{x}, \mathbf{W}_1)}{\psi_{s_k}(\mathbf{x}, \mathbf{w}_1^k)}\right], \tag{12}$$

$$\frac{\partial}{\partial \mathbf{w}_2} F_\theta(\mathbf{x}) = \frac{1}{\sqrt{2d_1}} \sum_{\mathbf{s} \in \{-1,1\}^{d_1}} \mathbf{s} \, \text{erf}'\left(\frac{\mathbf{w}_2 \cdot \mathbf{s}}{\sqrt{2d_1}}\right) \Psi_{\mathbf{s}}(\mathbf{x}, \mathbf{W}_1), \quad \text{with } \text{erf}'(x) := \frac{2}{\sqrt{\pi}}e^{-x^2}. \tag{13}$$
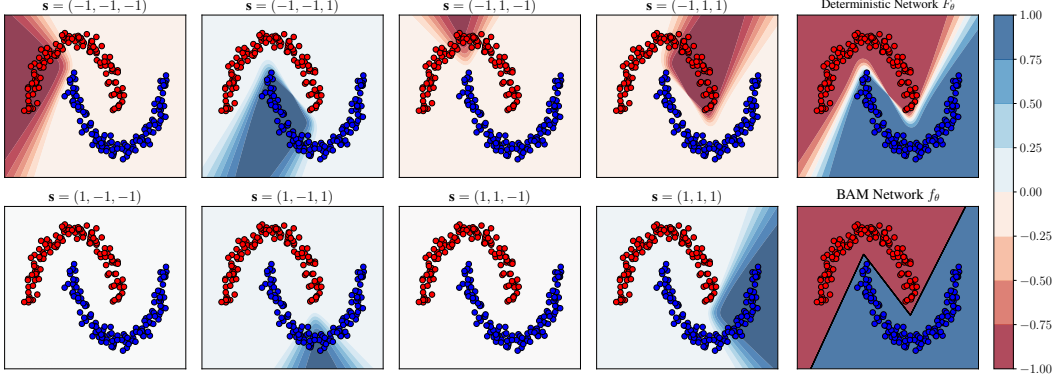
4

Figure 1: Illustration of the proposed method for a one hidden layer network of size $d_1{=}3$, interpreted as a majority vote over 8 binary representations $\mathbf{s} \in \{-1, 1\}^3$. For each $\mathbf{s}$, a plot shows the values of $F_{\mathbf{w}_2}(\mathbf{s})\Psi_{\mathbf{s}}(\mathbf{x}, \mathbf{W}_1)$. The sum of these values gives the deterministic network output $F_\theta(\mathbf{x})$ (see Eq. 9). We also plot the BAM network output $f_\theta(\mathbf{x})$ for the same parameters θ (see Eq. 6).

Note that this is an exact computation. A salient fact is that even though we work on non-differentiable BAM networks, we get a structure trainable by (stochastic) gradient descent by aggregating networks.

Majority vote of learned representations. Note that $\Psi_{\mathbf{s}}$ (Equation 10) defines a distribution on $\mathbf{s}$. Indeed, $\sum_{\mathbf{s}} \Psi_{\mathbf{s}}(\mathbf{x}, \mathbf{W}_1){=}1$, as $\Psi_{\mathbf{s}}(\mathbf{x}, \mathbf{W}_1) + \Psi_{\bar{\mathbf{s}}}(\mathbf{x}, \mathbf{W}_1) = 2^{-d_1}$ for every $\bar{\mathbf{s}} = -\mathbf{s}$. Thus, by Equation (9) we can interpret F_θ akin to a majority vote predictor, which performs a convex combination of a linear predictor outputs $F_{\mathbf{w}_2}(\mathbf{s}) := \mathrm{erf}(\mathbf{w}_2 \cdot \mathbf{s}/\sqrt{2d_1})$. The vote aggregates the predictions on the 2^{d_1} possible binary representations. Thus, the algorithm does not learn the representations *per se*, but rather the weights $\Psi_{\mathbf{s}}(\mathbf{x}, \mathbf{W}_1)$ associated to every $\mathbf{s}$ given an input $\mathbf{x}$, as illustrated by Figure 1.

3.2 Stochastic approximation

Since $\Psi_{\mathbf{s}}$ (Equation 10) defines a distribution, we can interpret the function value as the probability of mapping input $\mathbf{x}$ into the hidden representation $\mathbf{s}$ given the parameters $\mathbf{W}_1$. Using a different formalism, we could write $\Pr(\mathbf{s}|\mathbf{x}, \mathbf{W}_1) = \Psi_{\mathbf{s}}(\mathbf{x}, \mathbf{W}_1)$. This viewpoint suggests a sampling scheme to approximate both the predictor output (Equation 9) and the partial derivatives (Equations 12 and 13). We avoid computing the 2^{d_1} terms by resorting to a Monte Carlo approximation of the sum. Given an input $\mathbf{x}$ and a sampling size T, the procedure goes as follows.

Prediction. We generate T random binary vectors $Z{=}\{\mathbf{s}^t\}_{t=1}^{T}$ according to the $\Psi_{\mathbf{s}}(\mathbf{x}, \mathbf{W}_1)$-distribution. This can be done by uniformly sampling $z_i^t{\in}[0, 1]$, and setting $s_i^t{=}\mathrm{sgn}(\psi_1(\mathbf{x}, \mathbf{w}_1^i){-}z_i^t)$. A stochastic approximation of $F_\theta(\mathbf{x})$ is given by $\widehat{F}_\theta(Z) := \frac{1}{T}\sum_{t=1}^{T} \mathrm{erf}\left(\frac{\mathbf{w}_2 \cdot \mathbf{s}^t}{\sqrt{2d_1}}\right)$.

Derivatives. Note that for a given sample $\{\mathbf{s}^t\}_{t=1}^{T}$, the approximate derivatives according to $\mathbf{w}_2$ (Equation 15 below) can be computed numerically by the automatic differentiation mechanism of deep leaning frameworks while evaluating $\widehat{F}_\theta(Z)$ [*e.g.*, Paszke et al., 2017]. However, we need the following Equation (14) to approximate the gradient according to $\mathbf{W}_1$ because $\partial \widehat{F}_\theta(Z)/\partial \mathbf{w}_1^k = 0$.

$$\frac{\partial}{\partial \mathbf{w}_1^k} F_\theta(\mathbf{x}) \approx \frac{\mathbf{x}}{T\, 2^{\frac{3}{2}} \|\mathbf{x}\|} \mathrm{erf}'\left(\frac{\mathbf{w}_1^k \cdot \mathbf{x}}{\sqrt{2}\,\|\mathbf{x}\|}\right) \sum_{t=1}^{T} \frac{s_k^t}{\psi_{s_k^t}(\mathbf{x}, \mathbf{w}_1^k)} \mathrm{erf}\left(\frac{\mathbf{w}_2 \cdot \mathbf{s}^t}{\sqrt{2d_1}}\right), \tag{14}$$

$$\frac{\partial}{\partial \mathbf{w}_2} F_\theta(\mathbf{x}) \approx \frac{1}{T\sqrt{2d_1}} \sum_{t=1}^{T} \mathbf{s}^t\, \mathrm{erf}'\left(\frac{\mathbf{w}_2 \cdot \mathbf{s}^t}{\sqrt{2d_1}}\right) = \frac{\partial}{\partial \mathbf{w}_2} \widehat{F}_\theta(Z). \tag{15}$$

4 Generalization to multilayer networks

In the following, we extend the strategy introduced in Section 3 to BAM architectures with an arbitrary number of layers $L \in \mathbb{N}^*$ (Equation 1). An apparently straightforward approach to achieve this

generalization would have been to consider a Gaussian posterior distribution $\mathcal{N}(\theta, I_D)$ over the BAM family $\{f_{\tilde{\theta}} | \tilde{\theta} \in \mathbb{R}^D\}$. However, doing so leads to a deterministic network relying on undesirable sums of 2^D elements. Instead, we define a mapping $f_\theta \mapsto g_{\zeta(\theta)}$ which transforms the BAM network into a computation tree. An illustration of this mapping is provided in Appendix as Figure A.2.

BAM to tree architecture map. Given a BAM network f_θ of L layers with sizes $d_0, d_1, \ldots, d_L$ (reminder: $d_L = 1$), we obtain a *computation tree* by decoupling the neurons (*i.e.*, the computation graph nodes): the tree leaves contain $\prod_{k=1}^{L} d_k$ copies of each of the d_0 BAM input neurons, and the tree root node corresponds to the single BAM output neuron. Each input-output path of the original BAM network becomes a path of length L from one leaf to the tree root. Each tree edge has its own parameter (a real-valued scalar); the total number of edges/parameters is $D^\dagger := \sum_{k=0}^{L-1} d_k^\dagger$, with $d_k^\dagger := \prod_{i=k}^{L} d_i$. We define a tree parameters set η recursively according to the tree structure. From level k to $k+1$, the tree has $d_k^\dagger$ edges. That is, each node at level $k+1$ has its own parameters subtree $\eta^{k+1} := \{\eta_i^k\}_{i=0}^{d_k}$, where each η_i^k is either a weight vector containing the input edges parameters (by convention, $\eta_0^k \in \mathbb{R}^{d_{k-1}}$) or a parameter set (thus, $\eta_1^k, \ldots, \eta_{d_{k-1}}^k$ are themselves parameter subtrees). Hence, the *deepest* elements of the recursive parameters set η are weight vectors $\eta^1 \in \mathbb{R}^{d_0}$. Let us now define the output tree $g_\eta(\mathbf{x}) := g_L(\mathbf{x}, \eta)$ on an input $\mathbf{x} \in \mathbb{R}^{d_0}$ as a recursive function:

$$g_1(\mathbf{x}, \{\mathbf{w}\}) = \text{sgn}(\mathbf{w} \cdot \mathbf{x}),$$

$$g_{k+1}(\mathbf{x}, \underbrace{\{\mathbf{w}, \eta_1, \ldots, \eta_{d_k}\}}_{\eta}) = \text{sgn}\Big(\mathbf{w} \cdot \underbrace{(g_k(\mathbf{x}, \eta_1), \ldots, g_k(\mathbf{x}, \eta_{d_k}))}_{\mathbf{g}_k(\mathbf{x}, \eta)}\Big) \text{ for } k = 1, \ldots, L-1.$$

BAM to tree parameters map. Given BAM parameters θ, we denote $\theta_{1:k} := \text{vec}(\{\mathbf{W}_k\}_{i=1}^k)$. The mapping from θ into the corresponding (recursive) tree parameters set is $\zeta(\theta) = \{\mathbf{w}_L, \zeta_1(\theta_{1:L-1}), \ldots, \zeta_{d_{L-1}}(\theta_{1:L-1})\}$, such that $\zeta_i(\theta_{1:k}) = \{\mathbf{w}_k^i, \zeta_1(\theta_{1:k-1}), \ldots, \zeta_{d_{k-1}}(\theta_{1:k-1})\}$, and $\zeta_i(\theta_{1:1}) = \{\mathbf{w}_1^i\}$. Note that the parameters tree obtained by the transformation $\zeta(\theta)$ is highly redundant, as each weight vector $\mathbf{w}_k^i$ (the ith line of the $\mathbf{W}_k$ matrix from θ) is replicated $d_{k+1}^\dagger$ times. This construction is such that $f_\theta(\mathbf{x}) = g_{\zeta(\theta)}(\mathbf{x})$ for all $\mathbf{x} \in \mathbb{R}^{d_0}$.

Deterministic network With a slight abuse of notation, we let $\tilde{\eta} \sim Q_\eta := \mathcal{N}(\eta, I_{D^\dagger})$ denote a parameter tree of the same structure as η, where every weight is sampled *iid* from a normal distribution. We denote $G_\theta(\mathbf{x}) := \mathbf{E}_{\tilde{\eta} \sim Q_{\zeta(\theta)}} g_{\tilde{\eta}}(\mathbf{x})$, and we compute the output value of this predictor recursively. In the following, we denote $G_{\theta_{1:k+1}}^{(j)}(\mathbf{x})$ the function returning the jth neuron value of the layer $k+1$. Hence, the output of this network is $G_\theta(\mathbf{x}) = G_{\theta_{1:L}}^{(1)}(\mathbf{x})$. As such,

$$G_{\theta_{1:1}}^{(j)}(\mathbf{x}) = \int_{\mathbb{R}^{d_0}} Q_{\mathbf{w}_1^j}(\mathbf{v}) \text{sgn}(\mathbf{v} \cdot \mathbf{x}) d\mathbf{v} = \text{erf}\left(\frac{\mathbf{w}_1^j \cdot \mathbf{x}}{\sqrt{2}\|\mathbf{x}\|}\right),$$

$$G_{\theta_{1:k+1}}^{(j)}(\mathbf{x}) = \sum_{\mathbf{s} \in \{-1,1\}^{d_k}} \text{erf}\left(\frac{\mathbf{w}_{k+1}^j \cdot \mathbf{s}}{\sqrt{2d_k}}\right) \Psi_\mathbf{s}^k(\mathbf{x}, \theta), \text{ with } \Psi_\mathbf{s}^k(\mathbf{x}, \theta) = \prod_{i=1}^{d_k} \left(\frac{1}{2} + \frac{1}{2} s_i \times G_{\theta_{1:k}}^{(i)}(\mathbf{x})\right). \quad (16)$$

The complete mathematical calculations leading to the above results are provided in Appendix A.2. The computation tree structure and the parameter mapping $\zeta(\theta)$ are crucial to obtain the recursive expression of Equation (16). However, note that this abstract mathematical structure is never manipulated explicitly. Instead, it allows computing each hidden layer vector $(G_{\theta_{1:k}}^{(j)}(\mathbf{x}))_{j=1}^{d_k}$ sequentially; a summation of 2^{d_k} terms is required for each layer $k = 1, \ldots, L-1$.

Stochastic approximation. Following the Section 3.2 sampling procedure trick for the one hidden layer network, we propose to perform a stochastic approximation of the network prediction output, by a Monte Carlo sampling for each layer. Likewise, we recover exact and approximate derivatives in a layer-by-layer scheme. The related equations are given in Appendix A.3.

5 PBGNet: PAC-Bayesian SGD learning of binary activated networks

We design an algorithm to learn the parameters $\theta \in \mathbb{R}^D$ of the predictor G_θ by minimizing a PAC-Bayesian upper bound on the generalization loss $\mathcal{L}_\mathcal{D}(G_\theta)$. We name our algorithm PBGNet

(**P**AC-Bayesian **B**inary **G**radient **Net**work), as it is a generalization of the PBGD (PAC-Bayesian Gradient Descent) learning algorithm for linear classifiers [Germain et al., 2009] to deep binary activated neural networks.

Kullback-Leibler regularization. The computation of a PAC-Bayesian bound value relies on two key elements: the empirical loss on the training set and the Kullback-Leibler divergence between the prior and the posterior. Sections 3 and 4 present exact computation and approximation schemes for the empirical loss $\widehat{\mathcal{L}}_S(G_\theta)$ (which is equal to $\widehat{\mathcal{L}}_S(F_\theta)$ when $L{=}2$). Equation (17) introduces the KL-divergence associated to the parameter maps of Section 4. We use the shortcut notation $\mathcal{K}(\theta, \mu)$ to refer to the divergence between two multivariate Gaussians of $D^\dagger$ dimensions, corresponding to learned parameters $\theta = \text{vec}\left(\{\mathbf{W}_k\}_{k=1}^L\right)$ and prior parameters $\mu = \text{vec}\left(\{\mathbf{U}_k\}_{k=1}^L\right)$.

$$\mathcal{K}(\theta, \mu) \coloneqq \text{KL}\Big(Q_{\zeta(\theta)} \,\|\, P_{\zeta(\mu)}\Big) = \frac{1}{2}\left(\|\mathbf{w}_L - \mathbf{u}_L\|^2 + \sum_{i=1}^{L-1} d_{k+1}^\dagger \|\mathbf{W}_i - \mathbf{U}_i\|_F^2\right), \quad (17)$$

where the factors $d_{k+1}^\dagger = \prod_{i=k+1}^L d_i$ are due to the redundancy introduced by transformation $\zeta(\cdot)$. This has the effect of penalizing more the weights on the first layers. It might have a considerable influence on the bound value for very deep networks. On the other hand, we observe that this is consistent with the *fine-tuning* practice performed when training deep neural networks for a transfer learning task: prior parameters are learned on a first dataset, and the posterior weights are learned by adjusting the last layer weights on a second dataset [see Bengio, 2009, Yosinski et al., 2014].

Bound minimization. PBGNet minimizes the bound of Theorem 1 (rephrased as Equation 18). However, this is done indirectly by minimizing a variation on Theorem 2 and used in a deep learning context by Zhou et al. [2019] (Equation 19). Theorem 3 links both results (proof in Appendix A.4).

Theorem 3. *Given prior parameters $\mu \in \mathbb{R}^D$, with probability at least $1 - \delta$ over $S \sim \mathcal{D}^{\otimes n}$, we have for all θ on $\mathbb{R}^D$:*

$$\mathcal{L}_{\mathcal{D}}(G_\theta) \leq \underset{0<p<1}{\text{argsup}}\left\{\text{kl}(\widehat{\mathcal{L}}_S(G_\theta)\|p) \leq \frac{1}{n}[\mathcal{K}(\theta,\mu) + \ln\tfrac{2\sqrt{n}}{\delta}]\right\} \quad (18)$$

$$= \inf_{C>0}\left\{\tfrac{1}{1-e^{-C}}\left(1 - \exp\left(-C\,\widehat{\mathcal{L}}_S(G_\theta) - \tfrac{1}{n}[\mathcal{K}(\theta,\mu) + \ln\tfrac{2\sqrt{n}}{\delta}]\right)\right)\right\}. \quad (19)$$

We use stochastic gradient descent (SGD) as the optimization procedure to minimize Equation (19) with respect to θ and C. It optimizes the same trade-off as in Equation (5), but choosing the C value which minimizes the bound. The originality of our SGD approach is that not only do we induce gradient randomness by selecting *mini-batches* among the training set S, we also approximate the loss gradient by sampling T elements for the combinatorial sum at each layer. Our experiments show that, for some learning problems, reducing the sample size of the Monte Carlo approximation can be beneficial to the stochastic gradient descent. Thus the sample size value T has an influence on the cost function space exploration during the training procedure (see Figure 5 in the Appendix). Hence, we consider T as a PBGNet hyperparameter.

6 Numerical experiments

Experiments were conducted on six binary classification datasets, described in Appendix B.

Learning algorithms. In order to get insights on the trade-offs promoted by the PAC-Bayes bound minimization, we compared PBGNet to variants focusing on empirical loss minimization. We train the models using multiple network architectures (depth and layer size) and hyperparameter choices. The objective is to evaluate the efficiency of our PAC-Bayesian framework both as a learning algorithm design tool and a model selection criterion. For all methods, the network parameters are trained using the Adam optimizer [Kingma and Ba, 2015]. Early stopping is used to interrupt the training when the cost function value is not improved for 20 consecutive epochs. Network architectures explored range from 1 to 3 hidden layers (L) and a hidden size $h \in \{10, 50, 100\}$ ($d_i = h$ for $1 \leq i < L$). Unless otherwise specified, the same randomly initialized parameters are used as a prior in the bound and as a starting point for SGD optimization [as in Dziugaite and Roy, 2017]. Also, for all models except MLP, we select the binary activation sampling size T in a range going from 10 to 10000. More details about the experimental setting are given in Appendix B.

Table 1: Experiment results for the considered models on the binary classification datasets: error rates on the train and test sets (E_S and E_T), and generalization bounds on the linear loss $\mathcal{L}_\mathcal{D}$ (Bnd). The PAC-Bayesian bounds hold with probability 0.95. Bound values for PBGNet$_\ell$ are trivial and thus not reported. A visual representation of this table is presented in Appendix (Figure 4).

Dataset	MLP		PBGNet$_\ell$		PBGNet$_{\ell\text{-bnd}}$			PBGNet			PBGNet$_{pre}$		
	E_S	E_T	E_S	E_T	E_S	E_T	Bnd	E_S	E_T	Bnd	E_S	E_T	Bnd
ads	0.021	0.037	0.018	**0.032**	0.024	0.038	0.283	0.131	0.168	0.203	0.034	0.033	0.058
adult	0.128	0.149	0.136	**0.148**	0.158	0.154	0.227	0.157	0.154	0.219	0.153	0.151	0.165
mnist17	0.003	**0.004**	0.008	0.005	0.007	0.009	0.067	0.005	0.007	0.040	0.003	0.005	0.009
mnist49	0.002	**0.013**	0.003	0.018	0.034	0.039	0.153	0.035	0.040	0.139	0.018	0.021	0.030
mnist56	0.002	0.009	0.002	0.009	0.022	0.026	0.103	0.022	0.025	0.089	0.008	**0.008**	0.017
mnistLH	0.004	**0.017**	0.005	0.019	0.071	0.073	0.186	0.049	0.052	0.160	0.026	0.026	0.033

MLP. We compare to a standard network with `tanh` activation, as this activation resembles the `erf` function of PBGNet. We optimize the linear loss as the cost function and use 20% of training data as validation for hyperparameters selection. A weight decay parameter is selected between 0 and 10^{-4}. Using weight decay corresponds to adding an $L2$ regularizer to the cost function, but contrary to the regularizer of Equation (17) promoted by PBGNet, this regularization is uniform for all layers.

PBGNet$_\ell$. This variant minimizes the empirical loss $\widehat{\mathcal{L}}(G_\theta)$, with an $L2$ regularization term. The corresponding weight decay, as well as other hyperparameters, are selected using a validation set, exactly as the MLP does. The bound expression is not involved in the learning process and is computed on the model selected by the validation set technique.

PBGNet$_{\ell\text{-bnd}}$. Again, the empirical loss $\widehat{\mathcal{L}}(G_\theta)$ with an $L2$ regularization term is minimized. However, only the weight decay hyperparameter is selected on the validation set the other ones are selected by the bound. This method is motivated by an empirical observation: our PAC-Bayesian bound is a great model selection tool for most hyperparamers, except the weight decay term.

PBGNet. As described in Section 5, the generalization bound is directly optimized as the cost function during the learning procedure and used solely for hyperparameters selection: no validation set is needed and all training data S are exploited for learning.

PBGNet$_{pre}$. We also explore the possibility of using a part of the training data as a pre-training step. To do so, we split the training set into two halves. First, we minimize the empirical loss for a fixed number of 20 epochs on the first 50% of the training set. Then, we use the learned parameters as initialization and prior for PBGNet and learn on the second 50% of the training set.

Analysis. Results are summarized in Table 1, which highlights the strengths and weaknesses of the models. Both MLP and PBGNet$_\ell$ obtain competitive error scores but lack generalization guarantees. By introducing the bound value in the model selection process, even with the linear loss as the cost function, PBGNet$_{\ell\text{-bnd}}$ yields non-vacuous generalization bound values although with an increase in error scores. Using the bound expression for the cost function in PBGNet improves bound values while keeping similar performances. The Ads dataset is a remarkable exception where the small amount of training examples seems to radically constrain the network in the learning process as it hinders the KL divergence growth in the bound expression. With an informative prior from pre-training, PBGNet$_{pre}$ is able to recover competitive error scores while offering tight generalization guarantees. All selected hyperparameters are presented in Appendix (Table 4).

A notable observation is the impact of the bound exploitation for model selection on the train-test error gap. Indeed, PBGNet$_{\ell\text{-bnd}}$, PBGNet and PBGNet$_{pre}$ display test errors closer to their train errors, as compared to MLP and PBGNet$_\ell$. This behavior is more noticeable as the dataset size grows and suggests potential robustness to overfitting when the bound is used in the learning process.

7   Conclusion and perspectives

We made theoretical and algorithmic contributions towards a better understanding of generalization abilities of binary activated multilayer networks, using PAC-Bayes. Note that the computational complexity of a learning epoch of PBGNet is higher than the cost induced in *binary neural networks*

[Bengio, 2009, Hubara et al., 2016, 2017, Soudry et al., 2014]. Indeed, we focus on the optimization of the generalization guarantee more than computational complexity. Although we also propose a sampling scheme that considerably reduces the learning time required by our method, achieving a nontrivial tradeoff.

We intend to investigate how we could leverage the bound to learn suitable priors for PBGNet. Or equivalently, finding (from the bound point of view) the best network architecture. We also plan to extend our analysis to multiclass and multilabel prediction, and convolutional networks. We believe that this line of work is part of a necessary effort to gives rise to a better understanding of the behavior of deep neural networks.

Acknowledgments

References

Amiran Ambroladze, Emilio Parrado-Hernández, and John Shawe-Taylor. Tighter PAC-Bayes bounds. In *NIPS*, 2006.

Yoshua Bengio. Learning deep architectures for AI. *Foundations and Trends in Machine Learning*, 2 (1):1–127, 2009.

Olivier Catoni. A PAC-Bayesian approach to adaptive classification. *preprint*, 840, 2003.

Olivier Catoni. *Statistical learning theory and stochastic optimization: Ecole d'Eté de Probabilités de Saint-Flour XXXI-2001*. Springer, 2004.

Olivier Catoni. *PAC-Bayesian supervised classification: the thermodynamics of statistical learning*, volume 56. Inst. of Mathematical Statistic, 2007.

Corinna Cortes and Vladimir Vapnik. Support-vector networks. *Machine Learning*, 20(3), 1995.

Gintare Karolina Dziugaite and Daniel M. Roy. Computing nonvacuous generalization bounds for deep (stochastic) neural networks with many more parameters than training data. In *UAI*. AUAI Press, 2017.

Pascal Germain, Alexandre Lacasse, François Laviolette, and Mario Marchand. PAC-Bayesian learning of linear classifiers. In *ICML*, pages 353–360. ACM, 2009.

Ian Goodfellow, Yoshua Bengio, and Aaron Courville. *Deep Learning*. MIT Press, 2016. http://www.deeplearningbook.org.

Benjamin Guedj. A primer on PAC-Bayesian learning. *arXiv preprint arXiv:1901.05353*, 2019.

Itay Hubara, Matthieu Courbariaux, Daniel Soudry, Ran El-Yaniv, and Yoshua Bengio. Binarized neural networks. In *NIPS*, pages 4107–4115, 2016.

Itay Hubara, Matthieu Courbariaux, Daniel Soudry, Ran El-Yaniv, and Yoshua Bengio. Quantized neural networks: Training neural networks with low precision weights and activations. *JMLR*, 18 (1):6869–6898, 2017.

Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. In *ICLR*, 2015.

John Langford. Tutorial on practical prediction theory for classification. *JMLR*, 6, 2005.

John Langford and Rich Caruana. (Not) Bounding the True Error. In *NIPS*, pages 809–816. MIT Press, 2001.

John Langford and John Shawe-Taylor. PAC-Bayes & margins. In *NIPS*, 2002.

David McAllester. Some PAC-Bayesian theorems. *Machine Learning*, 37(3), 1999.

David McAllester. PAC-Bayesian stochastic model selection. *Machine Learning*, 51(1), 2003.

Behnam Neyshabur, Srinadh Bhojanapalli, and Nathan Srebro. A PAC-Bayesian approach to spectrally-normalized margin bounds for neural networks. In *ICLR*, 2018.

Emilio Parrado-Hernández, Amiran Ambroladze, John Shawe-Taylor, and Shiliang Sun. PAC-Bayes bounds with data dependent priors. *JMLR*, 13, 2012.

Adam Paszke, Sam Gross, Soumith Chintala, Gregory Chanan, Edward Yang, Zachary DeVito, Zeming Lin, Alban Desmaison, Luca Antiga, and Adam Lerer. Automatic differentiation in PyTorch. In *NIPS Autodiff Workshop*, 2017.

Matthias Seeger. PAC-Bayesian generalization bounds for gaussian processes. *JMLR*, 3, 2002.

John Shawe-Taylor and Robert C. Williamson. A PAC analysis of a Bayesian estimator. In *COLT*, 1997.

Daniel Soudry, Itay Hubara, and Ron Meir. Expectation backpropagation: Parameter-free training of multilayer neural networks with continuous or discrete weights. In *NIPS*, pages 963–971, 2014.

Leslie G Valiant. A theory of the learnable. In *Proceedings of the sixteenth annual ACM symposium on Theory of computing*, pages 436–445. ACM, 1984.

Jason Yosinski, Jeff Clune, Yoshua Bengio, and Hod Lipson. How transferable are features in deep neural networks? In *NIPS*, pages 3320–3328, 2014.

Wenda Zhou, Victor Veitch, Morgane Austern, Ryan P. Adams, and Peter Orbanz. Non-vacuous generalization bounds at the imagenet scale: a PAC-bayesian compression approach. In *ICLR*, 2019.

A Supplementary Material

A.1 From the *sign* activation to the *erf* function

For completeness, we present the detailed derivation of Equation (4). This results appear namely in Germain et al. [2009], Langford [2005], Langford and Shawe-Taylor [2002].

Given $\mathbf{x} \in \mathbb{R}^d$, we have

$$F_{\mathbf{w}}(\mathbf{x}) = \mathop{\mathbf{E}}_{\mathbf{v} \sim \mathcal{N}(\mathbf{w}, \mathrm{I})} \mathrm{sgn}(\mathbf{v} \cdot \mathbf{x})$$

$$= \int_{\mathbb{R}^d} \mathrm{sgn}(\mathbf{v} \cdot \mathbf{x}) \left(\frac{1}{\sqrt{2\pi}} \right)^d e^{-\frac{1}{2}\|\mathbf{v}-\mathbf{w}\|^2} d\mathbf{v}$$

$$= \int_{\mathbb{R}^d} (\mathbb{1}\left[\mathbf{v} \cdot \mathbf{x} > 0\right] - \mathbb{1}\left[\mathbf{v} \cdot \mathbf{x} < 0\right]) \left(\frac{1}{\sqrt{2\pi}} \right)^d e^{-\frac{1}{2}\|\mathbf{v}-\mathbf{w}\|^2} d\mathbf{v}$$

$$= \left(\frac{1}{\sqrt{2\pi}} \right)^d \int_{\mathbb{R}^d} \mathbb{1}\left[\mathbf{v} \cdot \mathbf{x} > 0\right] e^{-\frac{1}{2}\|\mathbf{v}-\mathbf{w}\|^2} d\mathbf{v} - \left(\frac{1}{\sqrt{2\pi}} \right)^d \int_{\mathbb{R}^d} \mathbb{1}\left[\mathbf{v} \cdot \mathbf{x} < 0\right] e^{-\frac{1}{2}\|\mathbf{v}-\mathbf{w}\|^2} d\mathbf{v}.$$

Without loss of generality, let us consider a vector basis where $\frac{\mathbf{x}}{\|\mathbf{x}\|}$ is the first coordinate. In this basis, the first elements of the vectors $\mathbf{v} = (v_1, v_2, \ldots, v_d)$ and $\mathbf{w} = (w_1, w_2, \ldots, w_d)$ are

$$v_1 = \frac{\mathbf{v} \cdot \mathbf{x}}{\|\mathbf{x}\|}, \qquad\qquad w_1 = \frac{\mathbf{w} \cdot \mathbf{x}}{\|\mathbf{x}\|}.$$

Hence, $\mathbf{v} \cdot \mathbf{x} = v_1 \cdot \|\mathbf{x}\|$ with $\|\mathbf{x}\| > 0$. Looking at the left side of the subtraction from the previous equation, we thus have

$$\left(\frac{1}{\sqrt{2\pi}} \right)^d \int_{\mathbb{R}^d} \mathbb{1}\left[\mathbf{v} \cdot \mathbf{x} > 0\right] e^{-\frac{1}{2}\|\mathbf{v}-\mathbf{w}\|^2} d\mathbf{v}$$

$$= \int_{\mathbb{R}} \mathbb{1}\left[v_1 > 0\right] \frac{1}{\sqrt{2\pi}} e^{-\frac{1}{2}(v_1-w_1)2} \left[\int_{\mathbb{R}^{d-1}} \left(\frac{1}{\sqrt{2\pi}} \right)^{d-1} e^{-\frac{1}{2}\|\mathbf{v}_{2:d}-\mathbf{w}_{2:d}\|^2} d\mathbf{v}_{2:d} \right] dv_1$$

$$= \frac{1}{\sqrt{2\pi}} \int_{-\infty}^{\infty} \mathbb{1}\left[t > -w_1\right] e^{-\frac{1}{2}t^2} dt,$$

with $t := v_1 - w_1$. Hence,

$$\mathop{\mathbf{E}}_{\mathbf{v} \sim \mathcal{N}(\mathbf{w}, \mathrm{I})} \mathrm{sgn}(\mathbf{v} \cdot \mathbf{x}) = \frac{1}{\sqrt{2\pi}} \int_{-\infty}^{\infty} \mathbb{1}\left[t > -w_1\right] e^{-\frac{1}{2}t^2} dt - \frac{1}{\sqrt{2\pi}} \int_{-\infty}^{\infty} \mathbb{1}\left[t < -w_1\right] e^{-\frac{1}{2}t^2} dt$$

$$= \frac{1}{\sqrt{2\pi}} \int_{-w_1}^{\infty} e^{-\frac{1}{2}t^2} dt - \frac{1}{\sqrt{2\pi}} \int_{-\infty}^{-w_1} e^{-\frac{1}{2}t^2} dt$$

$$= \frac{1}{2} + \frac{1}{\sqrt{2\pi}} \int_{0}^{w_1} e^{-\frac{1}{2}t^2} dt - \frac{1}{2} + \frac{1}{\sqrt{2\pi}} \int_{0}^{w_1} e^{-\frac{1}{2}t^2} dt$$

$$= \frac{\sqrt{2}}{\sqrt{\pi}} \int_{0}^{w_1} e^{-\frac{1}{2}t^2} dt$$

$$= \frac{2}{\sqrt{\pi}} \int_{0}^{\frac{w_1}{\sqrt{2}}} e^{-u^2} du \quad \text{with } u = \frac{t}{\sqrt{2}}$$

$$= \mathrm{erf}\left(\frac{w_1}{\sqrt{2}} \right)$$

$$= \mathrm{erf}\left(\frac{\mathbf{w} \cdot \mathbf{x}}{\sqrt{2}\,\|\mathbf{x}\|} \right),$$

where $\mathrm{erf}(\cdot)$ is the Gauss error function defined as $\mathrm{erf}(x) = \frac{2}{\sqrt{\pi}} \int_{0}^{x} e^{-t^2} dt$.
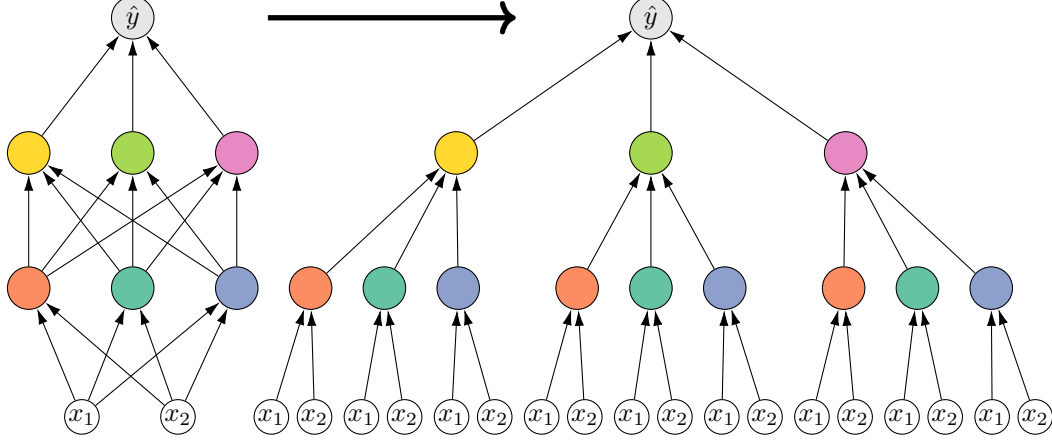
Figure 2: Illustration of the *BAM to tree architecture map* on a three layers network.

A.2 Prediction for the multilayer case

Details of the complete mathematical calculations leading to Equation (16) are presented below:

$$
\begin{aligned}
G^{(j)}_{\theta_{1:k+1}}(\mathbf{x}) &:= \int Q_{\zeta(\theta_{1:k+1})}(\tilde{\eta}) \, g_{k+1}(\mathbf{x}, \tilde{\eta}) \, d\tilde{\eta} \\
&= \int Q_{\zeta(\theta_{1:k})}(\tilde{\eta}_1) \dots \int Q_{\zeta(\theta_{1:k})}(\tilde{\eta}_{d_k}) \left(\int_{\mathbb{R}^{d^k}} Q_{\mathbf{w}_k^j}(\mathbf{v}) \mathrm{sgn}[\mathbf{v} \cdot \mathbf{g}_k(\mathbf{x}, \tilde{\eta})] d\mathbf{v} \right) d\tilde{\eta}_{d_k} \dots d\tilde{\eta}_1 \\
&= \int Q_{\zeta(\theta_{1:k})}(\tilde{\eta}_1) \dots \int Q_{\zeta(\theta_{1:k})}(\tilde{\eta}_{d_k}) \, \mathrm{erf}\left(\frac{\mathbf{w}_{k+1}^j \cdot \mathbf{g}_k(\mathbf{x}, \tilde{\eta})}{\sqrt{2}\|\mathbf{g}_k(\mathbf{x}, \tilde{\eta})\|} \right) d\tilde{\eta}_{d_k} \dots d\tilde{\eta}_1 \\
&= \sum_{\mathbf{s} \in \{-1,1\}^{d_k}} \mathrm{erf}\left(\frac{\mathbf{w}_{k+1}^j \cdot \mathbf{s}}{\sqrt{2d_k}} \right) \int Q_{\zeta(\theta_{1:k})}(\tilde{\eta}_1) \dots \int Q_{\zeta(\theta_{1:k})}(\tilde{\eta}_{d_k}) \mathbb{1}[\mathbf{s} = \mathbf{g}_k(\mathbf{x}, \tilde{\eta})] d\tilde{\eta}_{d_k} \cdots d\tilde{\eta}_1 \\
&= \sum_{\mathbf{s} \in \{-1,1\}^{d_k}} \mathrm{erf}\left(\frac{\mathbf{w}_{k+1}^j \cdot \mathbf{s}}{\sqrt{2d_k}} \right) \prod_{i=1}^{d_k} \int Q_{\zeta(\theta_{1:k})}(\tilde{\eta}_i) \mathbb{1}[s_i = g_k(\mathbf{x}, \tilde{\eta}_i)] d\tilde{\eta}_i \\
&= \sum_{\mathbf{s} \in \{-1,1\}^{d_k}} \mathrm{erf}\left(\frac{\mathbf{w}_{k+1}^j \cdot \mathbf{s}}{\sqrt{2d_k}} \right) \prod_{i=1}^{d_k} \int Q_{\zeta(\theta_{1:k})}(\tilde{\eta}_i) \left(\frac{1}{2} + \frac{s_i}{2} g_k(\mathbf{x}, \tilde{\eta}_i) \right) d\tilde{\eta}_i \\
&= \sum_{\mathbf{s} \in \{-1,1\}^{d_k}} \mathrm{erf}\left(\frac{\mathbf{w}_{k+1}^j \cdot \mathbf{s}}{\sqrt{2d_k}} \right) \prod_{i=1}^{d_k} \left(\frac{1}{2} + \frac{s_i}{2} \int Q_{\zeta(\theta_{1:k})}(\tilde{\eta}_i) g_k(\mathbf{x}, \tilde{\eta}_i) d\tilde{\eta}_i \right) \\
&= \sum_{\mathbf{s} \in \{-1,1\}^{d_k}} \mathrm{erf}\left(\frac{\mathbf{w}_{k+1}^j \cdot \mathbf{s}}{\sqrt{2d_k}} \right) \underbrace{\prod_{i=1}^{d_k} \left(\frac{1}{2} + \frac{1}{2} s_i \times G^{(i)}_{\theta_{1:k}}(\mathbf{x}) \right)}_{\Psi_{\mathbf{s}}^k(\mathbf{x}, \eta)}.
\end{aligned}
$$

Moreover,

$$
\Psi_{\mathbf{s}}^k(\mathbf{x}, \theta) = \prod_{i=1}^{d_k} \underbrace{\left(\frac{1}{2} + \frac{1}{2} s_i \times G^{(i)}_{\theta_{1:k}}(\mathbf{x}) \right)}_{\psi_{s_i}^k(\mathbf{x}, \theta)} .
$$

12

Base case:

$$G_{\theta_{1:1}}^{(j)}(\mathbf{x}) = \mathop{\mathbf{E}}_{\tilde{\eta} \sim \mathcal{N}(\zeta_j(\theta_{1:1}), I)} g_1(\mathbf{x}, \tilde{\eta})$$

$$= \int_{\mathbb{R}^{d_0}} Q_{\mathbf{w}_1^j}(\mathbf{v}) \mathrm{sgn}(\mathbf{v} \cdot \mathbf{x}) d\mathbf{v}$$

$$= \mathrm{erf}\left(\frac{\mathbf{w}_1^j \cdot \mathbf{x}}{\sqrt{2}\|\mathbf{x}\|}\right).$$

A.3   Derivatives of the multilayer case

We first aim at computing $\frac{\partial}{\partial \mathbf{w}_k} G_{\theta_{1:k}}(\mathbf{x})$.

Recall that $\mathbf{w}_k^j \in \{\mathbf{w}_k^1, \dots, \mathbf{w}_k^{d_k}\}$ is the j^{th} line of $\mathbf{W}_k$, that is the input weights of the corresponding hidden layer's neuron.

$$\frac{\partial}{\partial \mathbf{w}_{k+1}^j} G_{\theta_{1:k+1}}^{(j)}(\mathbf{x}) = \frac{\partial}{\partial \mathbf{w}_k^j} \sum_{\mathbf{s} \in \{-1,1\}^{d_k}} \mathrm{erf}\left(\frac{\mathbf{w}_{k+1}^j \cdot \mathbf{s}}{\sqrt{2d_k}}\right) \Psi_{\mathbf{s}}^k(\mathbf{x}, \theta)$$

$$= \sum_{\mathbf{s} \in \{-1,1\}^{d_k}} \frac{\mathbf{s}}{\sqrt{2d_k}} \mathrm{erf}'\left(\frac{\mathbf{w}_{k+1}^j \cdot \mathbf{s}}{\sqrt{2d_k}}\right) \Psi_{\mathbf{s}}^k(\mathbf{x}, \theta).$$

The base case of the recursion is

$$\frac{\partial}{\partial \mathbf{w}_1^j} G_{\theta_{1:1}}^{(j)}(\mathbf{x}) = \frac{\partial}{\partial \mathbf{w}_1^j} \mathrm{erf}\left(\frac{\mathbf{w}_1^j \cdot \mathbf{x}}{\sqrt{2}\|\mathbf{x}\|}\right)$$

$$= \frac{\mathbf{x}}{\sqrt{2}\|\mathbf{x}\|} \mathrm{erf}'\left(\frac{\mathbf{w}_1^j \cdot \mathbf{x}}{\sqrt{2}\|\mathbf{x}\|}\right).$$

In order to propagate the error through the layers, we also need to compute for $k > 1$:

$$\frac{\partial}{\partial G_{\theta_{1:k}}^{(l)}} G_{\theta_{1:k+1}}^{(j)} = \frac{\partial}{\partial G_{\theta_{1:k}}^{(l)}} \sum_{\mathbf{s} \in \{-1,1\}^{d_k}} \mathrm{erf}\left(\frac{\mathbf{w}_{k+1}^j \cdot \mathbf{s}}{\sqrt{2d_k}}\right) \prod_{i=1}^{d_k} \left(\frac{1}{2} + \frac{1}{2} s_i \times G_{\theta_{1:k}}^{(i)}\right)$$

$$= \sum_{\mathbf{s} \in \{-1,1\}^{d_k}} \mathrm{erf}\left(\frac{\mathbf{w}_{k+1}^j \cdot \mathbf{s}}{\sqrt{2d_k}}\right) \left[\frac{\Psi_{\mathbf{s}}^k(\mathbf{x}, \theta)}{\psi_{s_l^k}(\mathbf{x}, \theta)}\right] \frac{\partial}{\partial G_{\theta_{1:k}}^{(l)}} \psi_{s_l^k}(\mathbf{x}, \theta)$$

$$= \sum_{\mathbf{s} \in \{-1,1\}^{d_k}} \mathrm{erf}\left(\frac{\mathbf{w}_{k+1}^j \cdot \mathbf{s}}{\sqrt{2d_k}}\right) \left[\frac{s_l \Psi_{\mathbf{s}}^k(\mathbf{x}, \theta)}{2\psi_{s_l^k}(\mathbf{x}, \theta)}\right].$$

Thus, we can compute

$$\frac{\partial \widehat{\mathcal{L}}_S\left(G_{\theta_{1:k+1}}^{(j)}(\mathbf{x})\right)}{\partial \mathbf{w}_k^j} = \sum_l \frac{\partial \widehat{\mathcal{L}}_S\left(G_{\theta_{1:k+1}}^{(j)}(\mathbf{x})\right)}{\partial G_{\theta_{1:k+1}}^{(l)}} \frac{\partial G_{\theta_{1:k+1}}^{(l)}}{\partial G_{\theta_{1:k}}^{(j)}} \frac{\partial G_{\theta_{1:k}}^{(j)}}{\partial \mathbf{w}_k^j}.$$

A.4 Proof of Theorem 3

Lemma 4 (Germain et al. [2009], Proposition 2.1). *For any $0 < q \le p < 1$, we have*

$$\sup_{C>0}\left[\Delta(C, q, p)\right] = \mathrm{kl}(q\|p)\,,$$

with

$$\Delta(C, q, p) \coloneqq -\ln(1 - p(1 - e^{-C})) - Cq\,. \tag{20}$$

Proof. For $0 < q, p < 1$, $\Delta(C, q, p)$ is concave in C and the maximum is $c_0 = -\ln\left(\frac{qp-p}{qp-q}\right)$. Moreover, $\Delta(c_0, q, p) = \mathrm{kl}(q\|p)$. □

Theorem 3. *Given prior parameters $\mu \in \mathbb{R}^D$, with probability at least $1 - \delta$ over $S \sim \mathcal{D}^{\otimes n}$, we have for all θ on $\mathbb{R}^D$:*

$$\mathcal{L}_{\mathcal{D}}(G_\theta) \le \operatorname*{argsup}_{0<p<1}\left\{\mathrm{kl}(\widehat{\mathcal{L}}_S(G_\theta)\|p) \le \frac{1}{n}[\mathcal{K}(\theta, \mu) + \ln\tfrac{2\sqrt{n}}{\delta}]\right\}$$

$$= \inf_{C>0}\left\{\tfrac{1}{1-e^{-C}}\left(1 - \exp\left(-C\,\widehat{\mathcal{L}}_S(G_\theta) - \frac{1}{n}[\mathcal{K}(\theta, \mu) + \ln\tfrac{2\sqrt{n}}{\delta}]\right)\right)\right\}.$$

Proof. Let us define $p^* = \operatorname{argsup}_{0<p<1}\left\{\mathrm{kl}(\widehat{\mathcal{L}}_S(G_\theta)\|p) \le \frac{1}{n}[\mathcal{K}(\theta, \mu) + \ln\tfrac{2\sqrt{n}}{\delta}]\right\}.$

First, by a straightforward rewriting of Theorem 1 [Seeger, 2002], we have $\mathcal{L}_{\mathcal{D}}(G_\theta) \le p^*$.

Then, we want to show

$$p^* = \inf_{C>0}\left\{\tfrac{1}{1-e^{-C}}\left(1 - \exp\left(-C\,\widehat{\mathcal{L}}_S(G_\theta) - \frac{1}{n}[\mathcal{K}(\theta, \mu) + \ln\tfrac{2\sqrt{n}}{\delta}]\right)\right)\right\}. \tag{21}$$

By the definition of p^*, we have

$$\mathrm{kl}(\widehat{\mathcal{L}}_S(G_\theta)\|p^*) = \frac{1}{n}[\mathcal{K}(\theta, \mu) + \ln\tfrac{2\sqrt{n}}{\delta}],$$

and, by Lemma 4,

$$\sup_{C>0}\left[\Delta(C, \widehat{\mathcal{L}}_S(G_\theta), p^*)\right] = \frac{1}{n}[\mathcal{K}(\theta, \mu) + \ln\tfrac{2\sqrt{n}}{\delta}]\,, \tag{22}$$

$$\text{and } \forall C > 0 : \Delta(C, \widehat{\mathcal{L}}_S(G_\theta), p^*) \le \frac{1}{n}[\mathcal{K}(\theta, \mu) + \ln\tfrac{2\sqrt{n}}{\delta}]\,. \tag{23}$$

Let $C^* = \operatorname{argsup}_{C>0}\left[\Delta(C, \widehat{\mathcal{L}}_S(G_\theta), p^*)\right].$

By rearranging the terms of $\Delta(C^*, \widehat{\mathcal{L}}_S(G_\theta), p^*)$ (see Equation 20), we obtain, from Line (22),

$$p^* = \tfrac{1}{1-e^{-C^*}}\left(1 - \exp\left(-C^*\,\widehat{\mathcal{L}}_S(G_\theta) - \frac{1}{n}[\mathcal{K}(\theta, \mu) + \ln\tfrac{2\sqrt{n}}{\delta}]\right)\right), \tag{24}$$

and, from Line (23),

$$\forall C > 0 : p^* \le \tfrac{1}{1-e^{-C}}\left(1 - \exp\left(-C\,\widehat{\mathcal{L}}_S(G_\theta) - \frac{1}{n}[\mathcal{K}(\theta, \mu) + \ln\tfrac{2\sqrt{n}}{\delta}]\right)\right). \tag{25}$$

Thus, combining Lines (24) and (25), we finally prove the desired result of Equation (21). □

Table 2: Datasets overview.

Dataset	n_{train}	n_{test}	d
ads	2459	820	1554
adult	32561	16281	108
mnist17	11377	3793	784
mnist49	10336	3446	784
mnist56	9891	3298	784
mnistLH	52500	17500	784

Table 3: Models overview.

Model name	Cost function	Train split	Valid split	Model selection	Prior
MLP	linear loss, L2 regulatized	80%	20%	valid linear loss	-
PBGNet$_\ell$	linear loss, L2 regulatized	80%	20%	valid linear loss	random init
PBGNet$_{\ell\text{-bnd}}$	linear loss, L2 regulatized	80%	20%	hybrid (see B.2)	random init
PBGNet	PAC-Bayes bound	100 %	-	PAC-Bayes bound	random init
PBGNet$_{pre}$					
– pretrain	linear loss (20 epochs)	50%	-	-	random init
– final	PAC-Bayes bound	50%	-	PAC-Bayes bound	pretrain

B   Experiments

B.1   Datasets.

In Section 6 we use the datasets Ads (a small dataset related to advertisements on web pages), Adult (a low-dimensional task about predicting income from census data) and four binary variants of the MNIST handwritten digits:

ads http://archive.ics.uci.edu/ml/datasets/Internet+Advertisements
  The first 4 features which have missing values are removed.

adult https://archive.ics.uci.edu/ml/datasets/Adult

mnist http://yann.lecun.com/exdb/mnist/
  Binary classification tasks are compiled with the following pairs:

  - Digits pairs 1 vs. 7, 4 vs. 9 and 5 vs. 6.
  - Low digits vs. high digits ($\{0, 1, 2, 3, 4\}$ vs $\{5, 6, 7, 8, 9\}$) identified as *mnistLH*.

We split the datasets into training and testing sets with a 75/25 ratio except for adult which has a training/test split already computed. Table 2 presents an overview of the datasets statistics.

B.2   Learning algorithms details.

Table 3 summarizes the characteristics of the learning algorithms used in the experiments.

Hyperparameter choices.   We execute each learning algorithm for combination of hyperparameters selected among the following values.

  - Hidden layers $\in \{1, 2, 3\}$.
  - Hidden size $\in \{10, 50, 100\}$.
  - Sample size $\in \{10, 50, 100, 1000, 10000\}$.
  - Weight decay $\in \{0, 10^{-4}, 10^{-6}\}$.
  - Learning rate $\in \{0.1, 0.01, 0.001\}$.

Note that the sample size does not apply to MLP and weight decay is set to 0 for PBGNet and PBGNet$_{pre}$).

We report all the hyperparameters of selected models, for all learning algorithms and all datasets, in Table 4. The errors and bounds for these selected models are presented by Table 1 of the main paper. The same results are visually illustrated by Figure 4.

Hybrid model selection scheme. Of note, PBGNet$_{\ell\text{-bnd}}$ has a unique hyperparameters selection approach using a combination of the validation loss and the bound value. First all hyperparameters, except the weight decay, are selected in order to minimize the bound value. This includes choosing the best epoch from which loading the network weights. Thus, we obtain the best models according to the bound for each weight decay values considered. Then, the validation loss can be used to identify the best model between those, hence selecting the weight decay value.

Optimization. For all methods, the network parameters are trained using the Adam optimizer [Kingma and Ba, 2015]. for a maximum of 150 epochs on mini-batches of size 32 for the smaller datasets (Ads and MNIST digit pairs) and size 64 for bigger datasets (Adult and mnistLH). Initial learning rate is selected in $\{0.1, 0.01, 0.001\}$ and halved after each 5 consecutive epochs without a decrease in the cost function value. We empirically observe that the prediction accuracy of PBGNet is usually better when trained using Adam optimizer than with *plain* stochastic gradient descent, while both optimization algorithms give comparable results for our MLP model. The study of this phenomena is considered as an interesting future research direction.

Usually in deep learning framework training loops, the empirical loss of an epoch is computed as the averaged loss of each mini-batch. As the weights are updated after each mini-batch, the resulting epoch loss is only an approximation for the empirical loss of the final mini-batch weights. The linear loss being a significant element of the PAC-Bayesian bound expression, the approximation has a non-negligible impact over the corresponding bound value. One could obtain the accurate empirical loss for each epoch by assessing the network performance on the complete training data at the end of each epoch. We empirically evaluated that doing so leads to an increase of about a third of the computational cost per epoch for the inference computation. A practical alternative used in our experiments is to simply rely on the averaged empirical loss on the mini-batches in the bound expression for epoch-related actions: learning rate reduction, early stopping and best epoch selection.

Prediction. Once the best epoch is selected, we can afford to compute the correct empirical loss for those weights and use it to obtain the corresponding bound value. However, because PBGNet and its variants use a Monte Carlo approximation in the inference stage, the predicted output is not deterministic. Thus, to obtain the values reported in Table 1, we repeated the prediction over the training data 20 times for the empirical loss computation of the selected epoch. The inference repetition process was also carried out on the testing set, hence reported values E_S, E_T and Bnd of the results consist in the mean over 20 approximated predictions. The standard deviations are consistently below 0.001. If network prediction consistency is crucial, one can set a higher sample size during inference to decrease variability, but keep a smaller sample size during training to reduce computational costs.

Implementation details. We implemented PBGNet using `PyTorch` library [Paszke et al., 2017].

When computing the full combinatorial sum, a straightforward implementation is feasible, gradients being computed efficiently by the automatic differentiation mechanism. For speed purposes, we compute the sum as a matrix operation by loading all $\mathbf{s} \in \{-1, 1\}^{d_k}$ as an array. Thus we are mainly limited by memory usage on the GPU, a single hidden layer of hidden size 20 using approximately 10Gb of memory depending on the dataset and batch size used.

For the Monte Carlo approximation, we need to insert the gradient approximation in a flexible way into the derivative graph of the automatic differentiation mechanism. Therefore, we implemented each layer as a function of the weights and the output of the previous layer, with explicit forward and backward expression. Thus the automatic differentiation mechanism is able to accurately propagate the gradient through our approximated layers, and also combine gradient from other sources towards the weights (for example the gradient from the KL computation when optimizing with the bound as the cost function).

Experiments were performed on NVIDIA GPUs (Tesla V100, Titan Xp, GeForce GTX 1080 Ti).

B.3   Additional results

Figure 3 reproduces the experiment presented by Figure 1 with another toy dataset. Figure 5 studies the effect of the sampling size T on the stochastic gradient descent procedure. See both figures captions for details.
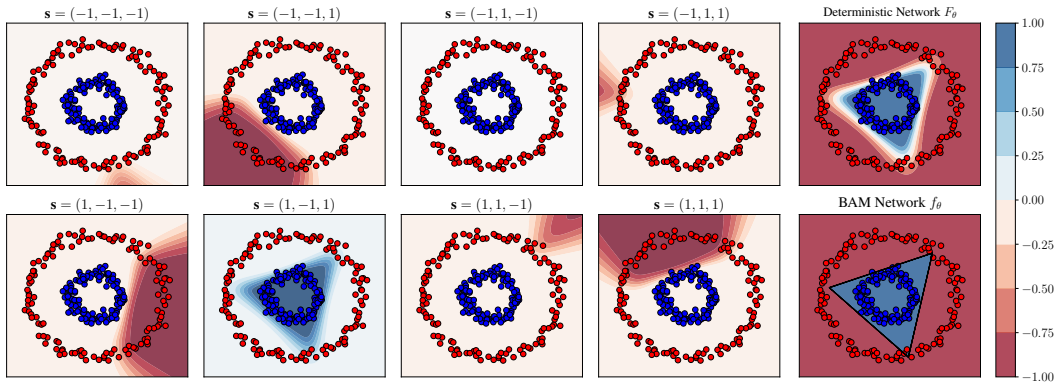
Figure 3: Illustration of the proposed method in Section 3 for a one hidden layer network of size $d_1 = 3$, interpreted as a majority vote over 8 binary representations $\mathbf{s} \in \{-1, 1\}^3$. For each $\mathbf{s}$, a plot shows the values of $F_{\mathbf{w}_2}(\mathbf{s})\Psi_{\mathbf{s}}(\mathbf{x}, \mathbf{W}_1)$. The sum of these values gives the deterministic network output $F_\theta(\mathbf{x})$ (see Equation 9). We also show the BAM network output $f_\theta(\mathbf{x})$ for the same parameters θ (see Equation 6).
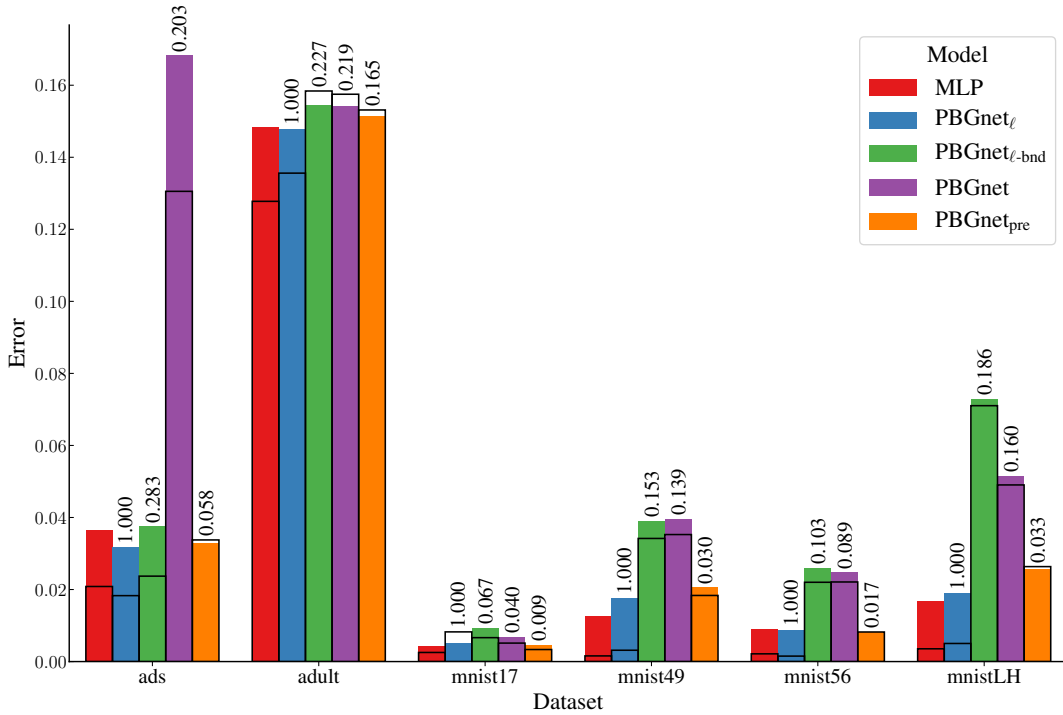

Figure 4: Visualization of experiment results for the models on the binary classification datasets. The colored bars display the test error while the black outlined bars exhibit the train error. The PAC-Bayesian generalization bounds are identified on the top of the bars and hold with probability 0.95.

Table 4: Selected models overview.

| Dataset	Model	Hid. layers	Hid. size	T	WD	C	KL	LR	Best epoch
ads	MLP	3	10	-	10^{-4}	-	-	0.01	49
	PBGNet$_\ell$	1	10	100	10^{-6}	12.25	15727	0.1	10
	PBGNet$_{\ell\text{-bnd}}$	1	10	10000	10^{-4}	2.08	360	0.01	66
	PBGNet	3	10	10000	-	0.47	27	0.01	49
	PBGNet$_{\text{pre}}$	2	10	1000	-	0.56	0.09	0.1	86
adult	MLP	2	10	-	10^{-6}	-	-	0.01	47
	PBGNet$_\ell$	3	50	10000	10^{-6}	19.99	6253241	0.01	66
	PBGNet$_{\ell\text{-bnd}}$	1	10	1000	10^{-4}	0.35	232	0.1	66
	PBGNet	1	10	10000	-	0.31	218	0.1	95
	PBGNet$_{\text{pre}}$	3	10	10000	-	0.09	0.07	0.01	80
mnist17	MLP	3	10	-	10^{-4}	-	-	0.1	22
	PBGNet$_\ell$	3	10	50	0	19.99	3522513	0.1	8
	PBGNet$_{\ell\text{-bnd}}$	1	10	10	10^{-6}	1.73	311	0.001	14
	PBGNet	1	10	10000	-	1.33	164	0.1	115
	PBGNet$_{\text{pre}}$	1	50	10000	-	0.73	0.30	0.1	100
mnist49	MLP	2	100	-	10^{-6}	-	-	0.001	46
	PBGNet$_\ell$	3	10	10000	0	19.99	541252	0.01	80
	PBGNet$_{\ell\text{-bnd}}$	1	10	10000	10^{-4}	1.03	335	0.1	27
	PBGNet	1	10	10000	-	0.90	305	0.1	91
	PBGNet$_{\text{pre}}$	1	50	1000	-	0.40	1.00	0.1	100
mnist56	MLP	2	50	-	0	-	-	0.001	23
	PBGNet$_\ell$	1	100	10000	0	19.99	532666	0.1	29
	PBGNet$_{\ell\text{-bnd}}$	1	10	100	10^{-4}	1.08	230	0.1	43
	PBGNet	1	10	1000	-	0.92	193	0.01	101
	PBGNet$_{\text{pre}}$	1	100	10000	-	0.54	0.61	0.1	74
mnistLH	MLP	2	100	-	10^{-6}	-	-	0.001	55
	PBGNet$_\ell$	3	100	10000	10^{-6}	19.99	63225412	0.01	38
	PBGNet$_{\ell\text{-bnd}}$	1	10	100	10^{-4}	0.50	628	0.1	40
	PBGNet	1	10	100	-	0.81	1541	0.1	102
	PBGNet$_{\text{pre}}$	2	100	10000	-	0.19	2.98	0.01	61
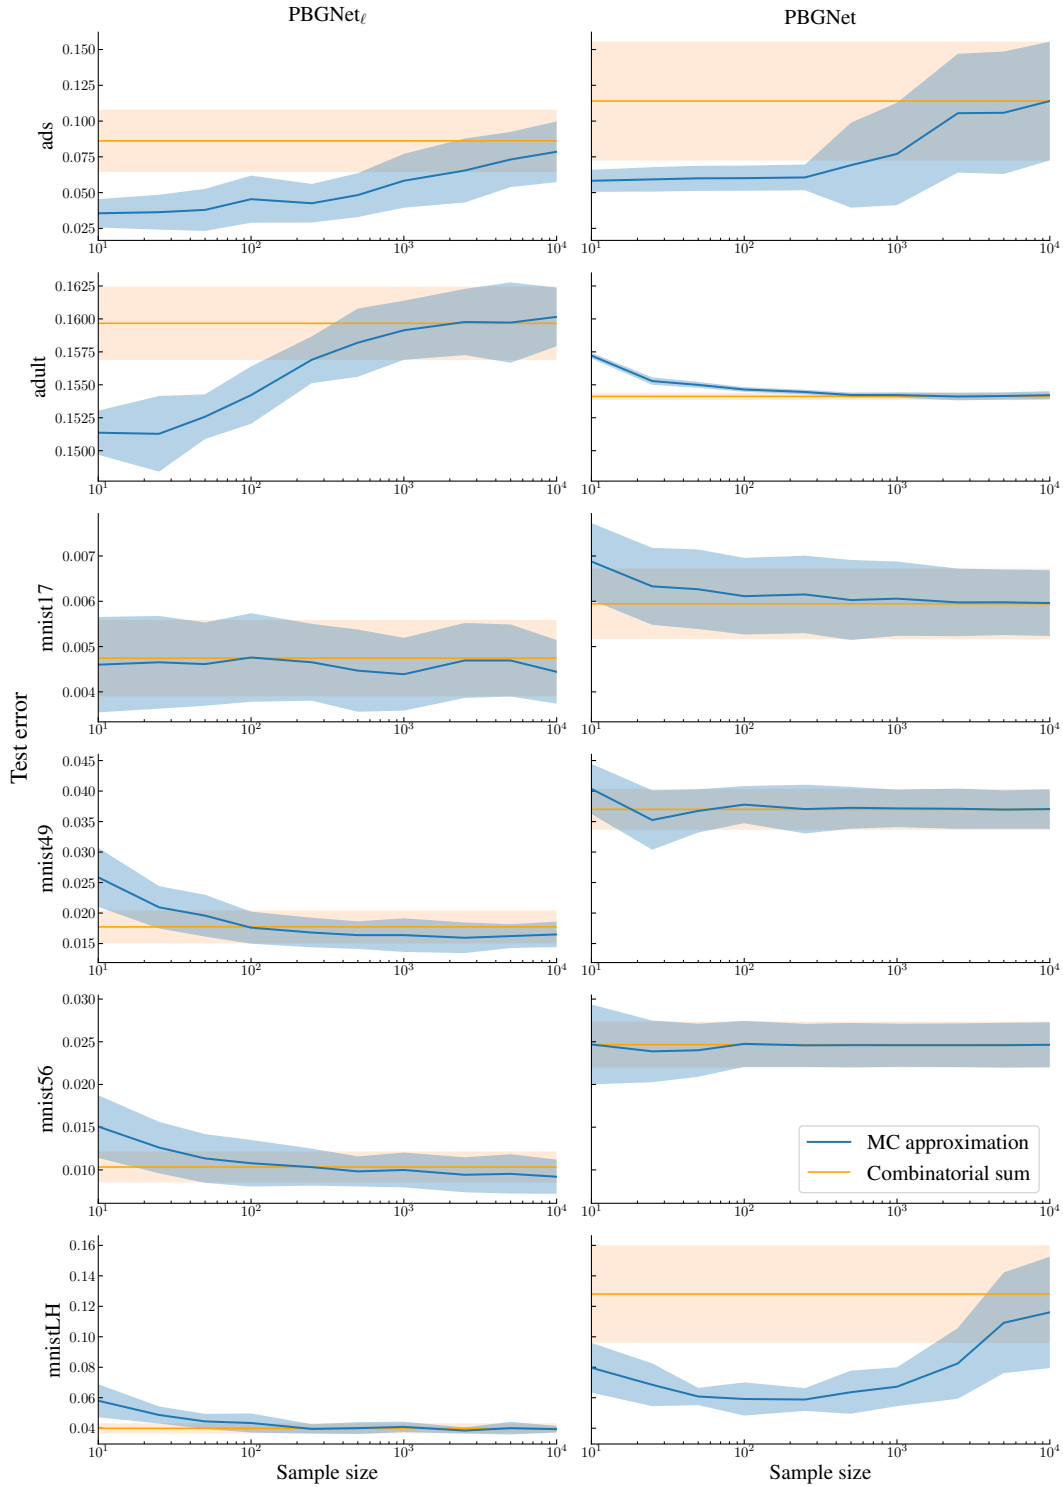
Figure 5: Impact of the sample size T on stochastic gradient descent solution test error for PBGNet$_\ell$ and PBGNet. Network parameters were fixed with a single hidden layer of size 10 and trained with initial learning rate of 0.1. For each sample size values and the combinatorial sum approach, 20 repetitions of the learning procedure with different random seeds were performed. The test error mean of the learned models is displayed by the bold line, with standard deviation shown with the shaded areas.