



HAL
open science

A path-norm toolkit for modern networks: consequences, promises and challenges

Antoine Gonon, Nicolas Brisebarre, Elisa Riccietti, Rémi Gribonval

► **To cite this version:**

Antoine Gonon, Nicolas Brisebarre, Elisa Riccietti, Rémi Gribonval. A path-norm toolkit for modern networks: consequences, promises and challenges. 2023. hal-04225201v1

HAL Id: hal-04225201

<https://hal.science/hal-04225201v1>

Preprint submitted on 2 Oct 2023 (v1), last revised 4 Dec 2024 (v5)

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

A PATH-NORM TOOLKIT FOR MODERN NETWORKS: CONSEQUENCES, PROMISES AND CHALLENGES

Antoine Gonon, Nicolas Brisebarre, Elisa Riccietti & Rémi Gribonval

Univ Lyon, EnsL, UCBL, CNRS, Inria, LIP, F-69342, LYON Cedex 07, France

ABSTRACT

This work introduces the first toolkit around path-norms that is fully able to encompass general DAG ReLU networks with biases, skip connections and max pooling. This toolkit notably allows us to establish generalization bounds for real modern neural networks that are not only the most widely applicable path-norm based ones, but also recover or beat the sharpest known bounds of this type. These extended path-norms further enjoy the usual benefits of path-norms: ease of computation, invariance under the symmetries of the network, and improved sharpness on feedforward networks compared to the product of operators' norms, another complexity measure most commonly used.

The versatility of the toolkit and its ease of implementation allow us to challenge the concrete promises of path-norm-based generalization bounds, by numerically evaluating the sharpest known bounds for ResNets on ImageNet.

1 INTRODUCTION

Developing a thorough understanding of theoretical properties of neural networks is key to achieve central objectives such as efficient and trustworthy training, robustness to adversarial attacks (e.g. via Lipschitz bounds), or statistical soundness guarantees (via so-called generalization bounds).

The so-called path-norm and path-embedding are promising concepts to theoretically analyze neural networks: the path-norm has been used to derive generalization guarantees (Neyshabur et al., 2015; Barron & Klusowski, 2019), and the path-embedding has lead for example to identifiability guarantees (Stock & Gribonval, 2022) and characterizations of properties of the dynamics of training algorithms (Marcotte et al., 2023).

However, the current definitions of the path-norm and of the path-embedding are currently severely limited: they only cover simple models unable to combine in a single framework pooling layers, skip connections, biases, or even multi-dimensional output (Neyshabur et al., 2015; Kawaguchi et al., 2017; Stock & Gribonval, 2022). Thus, the promises of existing theoretical guarantees based on these tools are currently out of reach as they cannot even be tested on standard modern networks.

Because of the current lack of versatility of these tools, known results have only been tested on toy examples. This prevents us from both understanding the reach of these tools and from diagnosing their strengths and weaknesses, which is necessary to either improve them in order to make them actually operational, if possible, or to identify without concession the gap between theory and practice, in particular for generalization bounds.

This work adresses the challenge of making these tools fully compatible with modern networks, and to concretely assess them on standard real-world examples. First, **it formalizes a definition of path-embedding (and path-norms) that is adapted to very generic ReLU networks**, covering any DAG architecture (in particular with skip connections), including in the presence of max/average pooling and/or biases. This covers a wide variety of modern networks (notably ResNets, VGGs, U-nets, ReLU MobileNets, Inception nets, Alexnet)¹, and recovers previously known definitions of these tools in simpler settings such as multilayer feedforward networks.

The immediate interests of these tools are: 1) *path-norms are easy to compute* on modern networks via a single forward-pass; 2) *path-norms are invariant under neuron permutations and parameter*

¹The conclusion discusses networks not covered by the framework and adaptations needed to cover them.

rescalings that leave the network invariant; and 3) the L^1 path-norm yields a Lipschitz bound of the network. These properties were known (but scattered in the literature) in the restricted case of feedforward ReLU networks (without max/average pooling or skip connections) primarily without biases (Neyshabur et al., 2015; Neyshabur, 2017; Furusho, 2020; Jiang et al., 2020; Dziugaite et al., 2020; Stock & Gribonval, 2022). They are generalized here for generic DAG ReLU networks with all the standard ingredients of modern networks.

Moreover, *path-norms tightly lower bound products of operator norms*, another complexity measure that does not enjoy the same invariances as path-norms, despite being widely used for Lipschitz bounds (e.g., to control adversarial robustness) (Neyshabur et al., 2018; Gonon et al., 2023) or generalization bounds (Neyshabur et al., 2015; Bartlett et al., 2017; Golowich et al., 2018). This bound, which was only known for feedforward ReLU networks without biases (Neyshabur et al., 2015), is again generalized here to generic DAG ReLU networks. This requires a proper adaptation for such networks that are not necessarily organized into layers with associated weight matrices.

Second, **this work also establishes a new generalization bound for modern ReLU networks based on their corresponding path-norm**. This bound covers arbitrary output dimension (while previous work focused on scalar dimension, see Table 1), generic DAG ReLU network architectures with max/average pooling, skip connections and biases. The achieved generalization bound recovers or beats the sharpest known ones of this type, that were so far only available in simpler restricted settings, see Table 1 for an overview. Among the technical ingredients used in the proof of this generalization bound, *the new contraction lemmas and the new peeling argument are among the main theoretical contributions of this work*. The first new contraction lemma extends the classical ones with scalar $t_i \in \mathbb{R}$, and contractions φ_i of the form $\mathbb{E}_\epsilon G(\sup_{t \in T} \sum_{i \in I} \epsilon_i \varphi_i(t_i)) \leq \mathbb{E}_\epsilon G(\sup_{t \in T} \sum_{i \in I} \epsilon_i t_i)$ (Ledoux & Talagrand, 1991, Theorem 4.12), with convex non-decreasing G , to situations where there are multiples independent copies indexed by $v \in V$ of the latter: $\mathbb{E}_\epsilon \max_{v \in V} G(\sup_{t \in T} \sum_{i \in I} \epsilon_{i,v} \varphi_i(t_i)) \leq \mathbb{E}_\epsilon \max_{v \in V} G(\sup_{t \in T} \sum_{i \in I} \epsilon_{i,v} t_i)$. The second new contraction lemma deals with vector-valued $t_i \in \mathbb{R}^W$, and $\varphi_i = \max$ to cope with max pooling neurons, and also multiple independent copies indexed by $v \in V$. The most related lemma we could find is the vector-valued ones in (Maurer, 2016) established with a different technique, and that holds only for $G = \text{id}$ with a single copy $|V| = 1$. The peeling argument allows to reduce the Rademacher complexity of the entire model to the one of the inputs by getting rid (peeling), one by one, of the neurons of the model. This is inspired by the peeling argument of (Golowich et al., 2018), which is however specific to feedforward ReLU networks with layer-wise constraints on the weights. Substantial additional ingredients are developed to handle *arbitrary DAG ReLU networks* (as there is no longer such a thing as a layer to peel), *with not only ReLU but also max and identity neurons* (where (Golowich et al., 2018) has only ReLU neurons), and *only a global constraint through the path-norm* (instead of layerwise constraints on operator norms). The analysis notably makes use of the rescaling invariance of the proposed generalized path-embedding.

The versatility of the proposed tools **enables us to compute for the first time generalization bounds based on path-norm on networks really used in practice**. This is the opportunity to assess the current state of the gap between theory and practice, and to diagnose possible room for improvements. As a concrete example, we demonstrate that on ResNet18 trained on ImageNet: 1) the proposed generalization bound can be numerically computed; 2) for a (dense) ResNet18 trained with standard tools, roughly *30 orders of magnitude* would need to be gained for this path-norm based bound to match practically observed generalization error; 3) the same bound evaluated on a *sparse* ResNet18 (trained with standard sparsification techniques) is decreased by up to 13 orders of magnitude. We conclude the paper by discussing promising leads to reduce this gap.

Paper structure. Section 2 introduces the ReLU networks being considered, and generalizes to this model the central definitions and results related to the path-embedding, the path-activations and the path-norm. Section 3 state a versatile generalization bound for such networks based on path-norm, and sketches its proof. Section 4 reports numerical experiments on ImageNet and ResNets. Related works are discussed along the way, and we refer to Appendix J for more details.

2 RELU MODEL AND PATH-EMBEDDING

Section 2.1 defines a general DAG ReLU model that covers modern architectures. Section 2.2 then introduces the so-called path-norm and extends related known results to this general model.

Table 1: Generalization bounds (up to universal multiplicative constants) for a ReLU network estimator in Θ learned from n iid training points when 1) the loss $\hat{y} \in (\mathbb{R}^{d_{\text{out}}}, \|\cdot\|_2) \rightarrow \ell(\hat{y}, y) \in \mathbb{R}$ is K -Lipschitz for every y , and 2) inputs are bounded in L^∞ -norm by $B \geq 1$. Here, $d_{\text{in}}/d_{\text{out}} =$ input/output dimension, $k =$ the maximum kernel size of the max pooling neurons, $M_d =$ matrix of layer d for a feedforward network (FFN), $D =$ depth. Note that r is more desirable than R since $r \leq R$ (Theorem B.1) and R can be arbitrarily large when $r = 0$ (Figure 2). This is because R decouples the layers without taking into account rescaling invariances.

	Architecture	Parameter set Θ	Generalization bound
(Kakade et al., 2008, Eq. (5))(Bach, Sec. 4.5.3)	linear regression (FFN, 0 hidden layer, $d_{\text{out}} = 1$)	$\ \theta\ _1 =$ $\ \Phi(\theta)\ _1 \leq r$	$\frac{KB}{\sqrt{n}} \times r \sqrt{\log(d_{\text{in}})}$
(E et al., 2022, Thm. 6) (Bach, 2017, Proposition 7)	one hidden layer, no biases, $d_{\text{out}} = 1$	$\ \Phi(\theta)\ _1 \leq r$	$\frac{KB}{\sqrt{n}} \times r \sqrt{\log(d_{\text{in}})}$
(Neyshabur et al., 2015, Corollary 7)	DAG, no biases, $d_{\text{out}} = 1$	$\ \Phi(\theta)\ _1 \leq r$	$\frac{KB}{\sqrt{n}} \times 2^D r \sqrt{\log(d_{\text{in}})}$
(Golowich et al., 2018, Theorem 3.2)	FFN, no biases, $d_{\text{out}} = 1$	$\prod_{d=1}^D \ M_d\ _{1,\infty} \leq R$	$\frac{KB}{\sqrt{n}} \times R \sqrt{D + \log(d_{\text{in}})}$
(Barron & Klusowski, 2019, Corollary 2)	FFN, no biases, $d_{\text{out}} = 1$	$\ \Phi(\theta)\ _1 \leq r$	$\frac{KB}{\sqrt{n}} \times r \sqrt{D + \log(d_{\text{in}})}$
Here, Theorem 3.1	DAG, with biases, arbitrary d_{out} , with ReLU, max pooling and identity neurons	$\ \Phi(\theta)\ _1 \leq r$	$\frac{KB}{\sqrt{n}} \times r \sqrt{D \log(k) + \log(d_{\text{in}} d_{\text{out}})}$

2.1 RELU MODEL THAT COVERS MODERN NETWORKS

The next definition introduces the model of ReLU neural networks being considered here.

Definition 2.1 (ReLU neural network). A ReLU neural network architecture is a DAG $G = (N, E, (\rho_v)_{v \in N})$ with edges E , and vertices N (called neurons) such that each neuron v has an attribute $\rho_v \in \{\text{id}, \text{max}, \text{ReLU}\}$ that corresponds to the activation function of v (identity, maximum or ReLU), with $\rho_v = \text{id}$ enforced whenever v has no successor. For a neuron v , the sets $\text{ant}(v), \text{suc}(v)$ of antecedents and successors of v are $\text{ant}(v) = \{u \in N, u \rightarrow v \in E\}, \text{suc}(v) = \{u \in N, v \rightarrow u \in E\}$. Neurons with no antecedents (resp. no successors) are called input (resp. output) neurons, and their set is denoted N_{in} (resp. N_{out}). Input and output dimensions are respectively $d_{\text{in}} = |N_{\text{in}}|$ and $d_{\text{out}} = |N_{\text{out}}|$. The set of neurons v with $\rho_v = \rho$ is denoted N_ρ . For $u \in N_{\text{max}}$, its kernel size is defined as being $|\text{ant}(u)|$.

Parameters associated with this architecture are vectors² $\theta \in \mathbb{R}^{E \cup N \setminus (N_{\text{in}} \cup N_{\text{max}})}$ (no biases on input neurons and max pooling neurons). We call bias and denote $b_v = \theta_v$ the coordinate associated with a neuron v , and denote $\theta^{u \rightarrow v}$ the weight associated with an edge $u \rightarrow v \in E$. We often denote $\theta^{\rightarrow v} = (\theta^{u \rightarrow v})_{u \in \text{ant}(v)}$ and $\theta^{v \rightarrow} = (\theta^{u \rightarrow v})_{u \in \text{suc}(v)}$.

In what follows, the symbol v can either denote a neuron $v \in N$ or the function associated to this neuron. The function $R_\theta^G : \mathbb{R}^{N_{\text{in}}} \rightarrow \mathbb{R}^{N_{\text{out}}}$ (simply denoted R_θ when G is clear from the context) realized by parameters θ is defined for every input $x \in \mathbb{R}^{N_{\text{in}}}$ as

$$R_\theta(x) := (v(\theta, x))_{v \in N_{\text{out}}},$$

where $v(\theta, x)$ is defined as $v(\theta, x) := x_v$ for an input neuron v , and defined by induction otherwise

$$v(\theta, x) := \begin{cases} \rho_v(b_v + \sum_{u \in \text{ant}(v)} \theta^{u \rightarrow v} u(\theta, x)) & \text{if } \rho_v = \text{ReLU} \text{ or } \rho_v = \text{id}, \\ \max_{u \in \text{ant}(v)} \theta^{u \rightarrow v} u(\theta, x) + b_v & \text{otherwise when } \rho_v = \text{max}. \end{cases}$$

Such a model indeed encompasses modern networks via the following implementations:

- *Max pooling*: set $\rho_v = \text{max}$, $b_v = 0$ and $\theta^{u \rightarrow v} = 1$ for every $u \in \text{ant}(v)$.
- *Average pooling*: set $\rho_v = \text{id}$, $b_v = 0$ and $\theta^{u \rightarrow v} = 1/|\text{ant}(v)|$ for every $u \in \text{ant}(v)$.
- *Batch normalization*: set $\rho_v = \text{id}$ and weights accordingly. Batch normalization layers only differ from standard affine layers by the way their parameters are updated during training.

²For an index set I , denote $\mathbb{R}^I = \{(\theta_i)_{i \in I}, \theta_i \in \mathbb{R}\}$.

- *Skip connections*: via the DAG structure, the outputs of any past layers can be added to the pre-activation of any neuron by adding connections from these layers to the given neuron.
- *Convolutional layers*: consider them as (doubly) circulant/Toeplitz fully connected layers.

2.2 PATH-EMBEDDING, PATH-ACTIVATIONS AND PATH-NORM

Given a general DAG ReLU network G as in Definition 2.1, it is possible to define a set of paths \mathcal{P}^G , a path-embedding Φ^G and path-activations \mathbf{A}^G , see Definition A.1 in the supplementary. The L^q path-norm is then $\|\Phi^G(\boldsymbol{\theta})\|_q$. Superscript G is omitted when obvious from context. The interest is that path-norm, which is easy to compute, can be interpreted as a Lipschitz bound of the network which is tighter than products of operator norms. We give here a high-level overview of definitions and properties, and refer to Appendix A for formal definitions, proofs and technical details. We highlight that the definitions and properties coincide with previously known ones in classical simpler settings.

Path-embedding and path-activations: fundamental properties. The main properties of the path-embedding Φ and the path-activations \mathbf{A} are: 1) for each parameter $\boldsymbol{\theta}$, the path-embedding $\Phi(\boldsymbol{\theta}) \in \mathbb{R}^{\mathcal{P}}$ is independent of the inputs x , and polynomial in the parameters $\boldsymbol{\theta}$ in a way that it is invariant under all neuron-wise rescaling symmetries³; 2) $\mathbf{A}(\boldsymbol{\theta}, x) \in \mathbb{R}^{\mathcal{P} \times (d_n+1)}$ takes a finite number of values and is piece-wise constant as a function of $(\boldsymbol{\theta}, x)$; and 3) denoting $\Phi^{\rightarrow v}$ and $\mathbf{A}^{\rightarrow v}$ to be the same objects but associated with the graph deduced from G by keeping only the largest subgraph of G with the same inputs as G and with single output v , then the output of every neuron v can be written as

$$v(\boldsymbol{\theta}, x) = \left\langle \Phi^{\rightarrow v}(\boldsymbol{\theta}), \mathbf{A}^{\rightarrow v}(\boldsymbol{\theta}, x) \begin{pmatrix} x \\ 1 \end{pmatrix} \right\rangle. \quad (1)$$

Compared to previous definitions given in simpler models (no pooling, no skip connections, no biases, one-dimensional output and/or layered network) (Kawaguchi et al., 2017; Stock & Gribonval, 2022), the main novelty is essentially to properly define the path-activations $\mathbf{A}(\boldsymbol{\theta}, x)$ in the presence of max pooling neurons: when going through a max pooling neuron, a path stays active only if the previous neuron of the path is the first in lexicographic order to realize the maximum.

Path-norm is easy to compute. It is mentioned in (Dziugaite et al., 2020, Appendix C.6.5) and (Jiang et al., 2020, Equation (44)) (without proof) that for ReLU *feedforward networks without biases*, the L^2 path-norm can be computed in a single forward pass with the formula: $\|\Phi^G(\boldsymbol{\theta})\|_2 = \|R_{|\boldsymbol{\theta}|^2}^G(\mathbf{1})\|_1$, where $|\boldsymbol{\theta}|^2$ is the vector $\boldsymbol{\theta}$ with $x \rightarrow x^2$ applied coordinate-wise (bias included) and where $\mathbf{1}$ is the constant input equal to one. This can be proved in a straightforward way, using Equation (1), and even extended to an arbitrary exponent $q \in [1, \infty]$: $\|\Phi^G(\boldsymbol{\theta})\|_q^q = \|R_{|\boldsymbol{\theta}|^q}^G(\mathbf{1})\|_1$. However, Appendix A shows that this formula is *false* when there are max pooling neurons, and that easy computation remains possible by first replacing the activation function of max pooling neurons with the identity before doing the forward pass. Average pooling neurons also need to be explicitly modeled as described after Definition 2.1.

L^1 path-norm yields a Lipschitz bound. Equation (1) is fundamental to understand the role of the path-norm. It shows that Φ contains information about the slopes of the function realized by the network on each of the regions where \mathbf{A} is constant. A formal result that goes in this direction is the Lipschitz bound $\|R_{\boldsymbol{\theta}}(x) - R_{\boldsymbol{\theta}}(x')\|_1 \leq \|\Phi(\boldsymbol{\theta})\|_1 \|x - x'\|_\infty$, already known in the case of ReLU feedforward networks (Neyshabur, 2017, before Section 3.4) (Furusho, 2020, Theorem 5), and generalized to the more general case of Definition 2.1 in Appendix A. This allows to leverage generic generalization bounds that apply to the set of all r -Lipschitz functions $f : [0, 1]^{d_n} \rightarrow [0, 1]$, however these bounds suffer from the curse of dimensionality (von Luxburg & Bousquet, 2004), unlike the bounds that we establish in Section 3.

Path-norms tightly lower bound products of operators' norms. For models of the form $R_{\boldsymbol{\theta}}(x) = M_D \text{ReLU}(M_{D-1} \dots \text{ReLU}(M_1 x))$, *i.e.* ReLU *feedforward* neural networks with matrices M_1, \dots, M_D and no biases, products such as $\prod_{d=1}^D \|M_d\|_{q, \infty}$, where $\|M\|_{q, \infty}$ is the maximum

³Because of positive-homogeneity of the considered activations functions, the realized function is preserved (Stock & Gribonval, 2022) when the incoming weights and the bias of a neuron are multiplied by $\lambda > 0$, while its outgoing weights are divided by λ . Path-norms inherit such symmetries and are further invariant to certain neuron permutations, typically within each layer in the case of feedforward networks.

L^q norm of a row of matrix M , can be used for $q = 1$, to bound the Lipschitz constant of the network (Neyshabur et al., 2018; Gonon et al., 2023) and to establish generalization bounds (Neyshabur et al., 2015; Bartlett et al., 2017; Golowich et al., 2018). So which one of path-norm and products of operator norms should be used? There are at least three reasons to consider the path-norm. **First**, it holds $\|\Phi(\theta)\|_q \leq \prod_{d=1}^D \|M_d\|_{q,\infty}$, with equality only if the parameters are properly rescaled. This is known for simple feedforward networks without biases (Neyshabur et al., 2015, Theorem 5). Appendix A generalizes it to DAGs as in Definition 2.1. The difficulty is to define the equivalent of the product of operators' norms with an arbitrary DAG and in the presence of biases. Apart from that, the proof is essentially the same as in (Neyshabur et al., 2015), with a similar rescaling that leads to equality of both measures, see Algorithm 1. **Second**, there are cases where the product of operators' norms is arbitrarily large while the path-norm is zero (see Figure 2 in Appendix B). Thus, it is *not desirable* to have a generalization bound that depends on this product of operators' norms since, compared to the path-norm, it fails to capture the complexity of the network end-to-end by *decoupling the layers of neurons* one from each other. **Third**, it has been empirically observed that products of operators' norms *negatively* correlate with the empirical generalization error while the path-norm *positively* correlates (Jiang et al., 2020, Table 2)(Dziugaite et al., 2020, Figure 1).

3 GENERALIZATION BOUND

The generalization bound of this section is based on path-norm for general DAG ReLU network. It encompasses modern networks and recovers or beats the sharpest known bounds of this type, and applies to the cross-entropy loss. The top-one accuracy loss is not directly covered, but can be controlled via a bound on the margin-loss, as detailed at the end of this section.

3.1 MAIN RESULT

To state the main result let us recall the definition of the generalization error.

Definition 3.1. Consider $d_{in}, d_{out} \in \mathbb{N}_{>0}$, $n \in \mathbb{N}_{>0}$ iid random variables $\mathbf{Z}_i = (\mathbf{X}_i, \mathbf{Y}_i) \in \mathbb{R}^{d_{in}} \times \mathbb{R}^{d_{out}}$. Consider an iid copy $\tilde{\mathbf{Z}}_1$ of \mathbf{Z}_1 . Consider a function $\ell : \mathbb{R}^{d_{out}} \times \mathbb{R}^{d_{out}} \rightarrow \mathbb{R}$. The ℓ -generalization error of any estimator $\hat{\theta}(\mathbf{Z}) \in \Theta$ is defined as:

$$\begin{aligned} \ell\text{-generalization error of } \hat{\theta}(\mathbf{Z}) := & \underbrace{\mathbb{E}_{\tilde{\mathbf{Z}}_1} \left(\ell \left(R_{\hat{\theta}(\mathbf{Z})}(\tilde{\mathbf{X}}_1), \tilde{\mathbf{Y}}_1 \right) \mid \mathbf{Z}_1, \dots, \mathbf{Z}_n \right)}_{\text{test error when trained on } \mathbf{Z}} \\ & - \underbrace{\frac{1}{n} \sum_{i=1}^n \ell \left(R_{\hat{\theta}(\mathbf{Z})}(\mathbf{X}_i), \mathbf{Y}_i \right)}_{\text{training error when trained on } \mathbf{Z}}. \end{aligned}$$

Theorem 3.1. Consider $d_{in}, d_{out} \in \mathbb{N}_{>0}$ and $n \in \mathbb{N}_{>0}$ iid random variables $\mathbf{Z}_i = (\mathbf{X}_i, \mathbf{Y}_i) \in \mathbb{R}^{d_{in}} \times \mathbb{R}^{d_{out}}$. Define $\sigma := \left(\mathbb{E}_{\mathbf{X}} \max \left(n, \max_{u=1, \dots, d_{in}} \sum_{i=1}^n (\mathbf{X}_i)_u^2 \right) \right)^{1/2}$. Consider a general DAG ReLU network as in Definition 2.1, with input dimension d_{in} and output dimension d_{out} . Denote D its depth (the maximal length of a path from an input to an output) and k its maximal kernel size (i.e. the maximum of $|\text{ant}(u)|$ over all max pooling neurons u), with $k = 1$ by convention when there is no max pooling neuron. Consider any loss function $\ell : \mathbb{R}^{d_{out}} \times \mathbb{R}^{d_{out}} \rightarrow \mathbb{R}$ such that

$$\ell(\hat{y}_1, y) - \ell(\hat{y}_2, y) \leq K \|\hat{y}_1 - \hat{y}_2\|_2, \quad \forall y, \hat{y}_1, \hat{y}_2 \in \text{support}(\mathbf{Y}_1). \quad (2)$$

for some $K > 0$. Consider a set of parameters Θ . Then⁴ for any estimator $\hat{\theta} : \mathbf{Z} \mapsto \hat{\theta}(\mathbf{Z}) \in \Theta$:

$$\mathbb{E}_{\mathbf{Z}} \ell\text{-generalization error of } \hat{\theta}(\mathbf{Z}) \leq \frac{4\sigma}{n} KC \sup_{\theta \in \Theta} \|\Phi(\theta)\|_1$$

with (log being the natural logarithm)

$$C := \left(D \log(5k) + \log \left(\frac{5}{2} (d_{in} + 1) d_{out} \right) \right)^{1/2}.$$

⁴Classical concentration results (Boucheron et al., 2013) can be used to deduce a bound that holds with high probability under additional mild assumptions on the loss.

Theorem 3.1 applies to the cross-entropy loss with $K = \sqrt{2}$ (see Appendix F) if the labels y are one-hot encodings⁵. A final softmax layer can be incorporated for free to the model by putting it in the loss. This does not change the bound since it is 1-Lipschitz with respect to the L^2 -norm (this is a simple consequence of the computations made in Appendix F).

On ImageNet, it holds $1/\sqrt{n} \leq \sigma/n \leq 2.6/\sqrt{n}$ (Section 4). This yields a bounds that decays in $\mathcal{O}(n^{-1/2})$ which is better than the generic $\mathcal{O}(n^{-1/d_{in}})$ generalization bound for Lipschitz functions (von Luxburg & Bousquet, 2004, Thm. 18) that suffer from the curse of dimensionality. Besides its wider range of applicability, this bounds also recovers or beats the sharpest known ones based on path-norm, see Table 1.

Sketch of proof for Theorem 3.1. The proof idea is explained below. Details are in Appendix E.

Already known ingredients Classical arguments (Shalev-Shwartz & Ben-David, 2014, Theorem 26.3)(Maurer, 2016), that are valid for any model, bound the expected generalization error by the Rademacher complexity of the model. It then remains to bound the latter, and this gets specific to neural networks. In the case of a feedforward ReLU neural network with no biases and scalar output (and no max pooling nor skip connections), (Golowich et al., 2018) proved that it is possible to bound this Rademacher complexity with no exponential factor in the depth, by peeling, one by one, each layer off the Rademacher complexity. To get more specific, for a class of functions \mathbf{F} and a function $\Psi : \mathbb{R} \rightarrow \mathbb{R}$, denote $\text{Rad} \circ \Psi(\mathbf{F}) = \mathbb{E}_\varepsilon \Psi(\sup_{f \in \mathbf{F}} \sum_{i=1}^n \varepsilon_i f(x_i))$ the Rademacher complexity of \mathbf{F} associated with n inputs x_i and Ψ . The goal for a generalization bound is to bound this in the case $\Psi(x) = \text{id}(x) = x$. In the specific case where \mathbf{F}_D is the class of functions corresponds to ReLU feedforward networks with depth D , assuming that some operator norm of each layer d is bounded by r_d , then (Golowich et al., 2018) basically guarantees $\text{Rad} \circ \Psi_\lambda(\mathbf{F}_D) \leq 2 \text{Rad} \circ \Psi_{\lambda r_D}(\mathbf{F}_{D-1})$ for every $\lambda > 0$, where $\Psi_\lambda(x) = \exp(\lambda x)$. Compared to previous works of (Golowich et al., 2018) that were directly working with $\Psi = \text{id}$ instead of Ψ_λ , the important point is that working with Ψ_λ gets the 2 outside of the exponential. Iterating over the depth D , optimizing over λ , and taking a logarithm at the end yields (by Jensen’s inequality) a bound on $\text{Rad} \circ \text{id}(\mathbf{F}_D)$ with a dependence on D that grows as $\sqrt{D \log(2)}$ instead of 2^D for previous approaches.

Novelties for general DAG ReLU networks Compared to the setup of (Golowich et al., 2018), there are at least three difficulties to do something similar here. First, *the neurons are not organized in layers* as the model can be an arbitrary DAG. So what should be peeled off one by one? Second, *the neurons are not necessarily ReLU neurons* as their activation function might be the identity (average pooling) or the maximum (max pooling). Finally, (Golowich et al., 2018) has a constraint on the weights of each layer, which makes it possible to pop out the constant r_d when layer d is peeled off. *Here, the only constraint is global*, since it constrains the paths of the network through $\|\Phi(\theta)\|_1 \leq r$. In particular, due to rescalings, the weights of a given neuron could be arbitrarily large or small under this constraint.

The first difficulty is primarily addressed using a new peeling lemma (Appendix D) that exploits a new contraction lemma (Appendix C).

The second difficulty is resolved by splitting the ReLU, max pooling and identity neurons in three groups before each peeling step. This has almost no impact on the final bound, as it only changes a $\log(2)$ factor in (Golowich et al., 2018) into a $\log(5)$ factor here.

Finally, the third obstacle is overcome by *rescaling the parameter* to normalize the vector of incoming weights of each neuron. This type of rescaling has also been used in (Neyshabur et al., 2015; Barron & Klusowski, 2019). \square

Remark 3.1 (Improved bound with assumptions on max pooling neurons). *In the specific case where max pooling neurons are grouped in layers, and when there are no skip connections going over these max pooling layers (satisfied by ResNets, not satisfied by U-nets), a sharpened peeling argument can yield the same bound but with C replaced by $C_{\text{sharpened}} = (D \log(3) + M \log(k) + \log((d_{in} + 1)d_{out}))^{1/2}$ with M being the number of max pooling layers (cf. Appendix D). The details are tedious so we only mention this result without proof. This ba-*

⁵A vector y is a one-hot encoding of a class c if $y = (\mathbb{1}_{c'=c})_{c' \in \{1, \dots, d_{out}\}}$.

sically replaces $\sqrt{D \log(5k)}$ by $\sqrt{D \log(3) + M \log(k)}$. For Resnet152, $k = 9$, $D = 152$ and $M = 1$, $\sqrt{D \log(5k)} \simeq 24$ while $\sqrt{D \log(3) + M \log(k)} \simeq 13$.

3.2 HOW TO DEAL WITH THE TOP-1 ACCURACY LOSS?

Theorem 3.1 does not apply to the top-1 accuracy loss as Equation (2) cannot be satisfied for any finite $K > 0$ in general (see Appendix G). It is still possible to bound the expected (test) top-1 accuracy by the so-called *margin loss* achieved at training (Bartlett et al., 2017, Lemma A.4). The margin-loss is a relaxed definition of the top-1 accuracy loss. A corollary of Theorem 3.1 is the next result proved in Appendix H.

Theorem 3.2 (Bound on the probability of misclassification). *Consider the setting of Theorem 3.1. Assume that the labels are indices $y \in \{1, \dots, d_{\text{out}}\}$. For any $\gamma > 0$, it holds*

$$\mathbb{P} \left(\arg \max_c R_{\theta}(\mathbf{X}_1)_c \neq \mathbf{Y}_1 \right) \leq \frac{1}{n} \sum_{i=1}^n \mathbb{1}_{(R_{\hat{\theta}(\mathbf{z})}(\mathbf{x}_i))_{\mathbf{Y}_i} \leq \gamma + \max_{c \neq \mathbf{Y}_i} (R_{\hat{\theta}(\mathbf{z})}(\mathbf{x}_i))_c} + \frac{8\sigma}{n} C \frac{\sup_{\theta} \|\Phi(\theta)\|_1}{\gamma}. \quad (3)$$

Note that the result is homogeneous: scaling both the outputs of the model and γ by the same scalar leaves the classifier and the corresponding bound unchanged.

4 EXPERIMENTS

Theorem 3.1 gives the first path-norm generalization bound that can be applied to modern networks (with max/average pooling, skip connections etc.). This bound is also the sharpest known bound of this type (Table 1). Since this bound is also easy to compute, the goal of this section is to numerically challenge for the first time the sharpest generalization bounds based on path-norm on modern networks. Note also that path-norms tightly lower bound products of operator norms (Appendix B) so that this also challenges the latter.

When would be the bound informative? For ResNets trained on ImageNet, the training error associated with cross-entropy is typically between 1 and 2, and the top-1 training error is typically less than 0.30. The same orders of magnitude apply to the empirical generalization error. To ensure that the test error (either for cross-entropy or top-1 accuracy) is of the same order as the training error, *the bound should basically be of order 1*.

For parameters θ learned from training data, Theorem 3.1 and Theorem 3.2 allow to bound the expected loss in terms of a performance measure (that depends on a free choice of $\gamma > 0$ for the top-1 accuracy) on training data plus a term bounded by $\frac{4\sigma}{n} C \times K \times \|\Phi(\theta)\|_1$. The Lipschitz constant K is $\sqrt{2}$ for cross-entropy, and $2/\gamma$ for the top-1 accuracy.

Evaluation of $\frac{4\sigma}{n} C$ for ResNets on ImageNet. We further bound σ/n by B/\sqrt{n} , where $B \simeq 2.6$ is the maximum L^∞ -norm of the images of ImageNet normalized for inference. We at most lose a factor B compared to the bound directly involving σ since it also holds $\sigma/n \geq 1/\sqrt{n}$ by definition of σ . We train on 99% of ImageNet so that $n = 1268355$. Moreover, $C = (D \log(5k) + \log(\frac{5}{2}(d_{\text{in}} + 1)d_{\text{out}}))^{1/2}$ with kernel size $k = 9$ for ResNets, $d_{\text{in}} = 224 \times 224 \times 3$ and $d_{\text{out}} = 1000$. The depth is $D = 3 + \# \text{ basic blocks} \times \# \text{ conv per basic block}$, with the different values available in Appendix I. The values for $4BC/\sqrt{n}$ are reported in Table 2. Given these results and the values of K , on ResNet18, *the bound would be informative only when $\|\Phi(\theta)\|_1 \lesssim 10$ or $\|\Phi(\theta)\|_1/\gamma \lesssim 10$ respectively for the cross-entropy and the top-1 accuracy.*

We now compute the path-norms of trained ResNets, both dense and sparse, using the simple formula proved in Theorem A.1 in appendix.

L^1 -path-norm of pretrained ResNets are 30 orders of magnitude too large. Table 3 shows that the L^1 path-norm is 30 orders of magnitude too large to make the bound informative for the cross-entropy loss. The choice of γ is discussed in Appendix I where we also show that there is no possible choice that leads to an informative bound for top-1 accuracy.

Table 2: Numerical evaluations on ResNets and ImageNet1k with 2 significant digits. Multiplying by the Lipschitz constant K of the loss and the path-norm gives the bound in Theorem 3.1. The second line reports the values when the analysis is sharpened for max pooling neurons, see Section 3.1.

ResNet	18	34	50	101	152
$\frac{4}{\sqrt{n}}CB =$	0.090	0.12	0.14	0.19	0.23
$\frac{4}{\sqrt{n}}C_{\text{sharpened}}B =$	0.061	0.072	0.082	0.11	0.13

Table 3: Path-norms of pretrained ResNets available on PyTorch, computed in float32.

ResNet	18	34	50	101	152
$\ \Phi(\theta)\ _1$	1.3×10^{30}	overflow	overflow	overflow	overflow
$\ \Phi(\theta)\ _2$	2.5×10^2	1.1×10^2	2.0×10^8	2.9×10^9	8.9×10^{10}
$\ \Phi(\theta)\ _4$	7.2×10^{-6}	4.9×10^{-6}	6.7×10^{-4}	3.0×10^{-4}	1.5×10^{-4}

Sparse ResNets can decrease the bounds by 13 orders of magnitude. We have just seen that pre-trained ResNets have very large L^1 path-norm. Does every network with a good test top-1 accuracy have a path-norm as large as this? Since any zero in the parameters θ leads to many coordinates of $\Phi(\theta)$ to be zero, we now investigate whether sparse versions of ResNet18 on ImageNet have a smaller path-norm. Sparse networks are obtained with iterative magnitude pruning plus rewinding, with hyperparameters similar to the one in (Frankle et al., 2021, Appendix A.3). Results show that the L^1 path-norm decreases from $\simeq 10^{30}$ for the dense network to $\simeq 10^{17}$ after 19 pruning iterations, basically losing between a half and one order of magnitude per pruning iteration. Moreover, the test top-1 accuracy is better than with the dense network for the first 11 pruning iterations, and after 19 iterations, the test top-1 accuracy is still way better than what would be obtained by guessing at random, so this is still a non-trivial matter to bound the generalization error for the last iteration. Details are in Appendix I. This shows that there are indeed practically trainable networks with much smaller L^1 path-norm that perform well. It remains open whether alternative training techniques, possibly with path-norm regularization, could lead to networks combining good performance and informative generalization bounds.

Additional observations: increasing the depth and the train size. In practice, increasing the size of the network (*i.e.* the number of parameters) or the number of training samples can improve generalization. We can, again, assess for the first time whether the bounds based on path-norms follows the same trend for standard modern networks. Table 3 shows that path-norms of pretrained ResNets available on PyTorch roughly increase with depth. This is complementary to (Dziugaite et al., 2020, Figure 1) where it is empirically observed on *simple feedforward* models that path-norm has difficulty to correlate positively with the generalization error when the depth evolves. For increasing training sizes, we did not observe a clear trend for the L^1 path-norm, which seems to mildly evolve with the number of epochs rather than with the train size, see Appendix I for details.

5 CONCLUSION

Contribution. To the best of our knowledge, this work is the first to introduce path-norm related tools for general DAG ReLU networks (with max/average pooling, skip connections), and Theorem 3.1 is the first generalization bound valid for such networks based on path-norm. This bound recovers or beats the sharpest known ones of the same type. Its ease of computation leads to the first experiments on modern networks that assess the promises of such approaches. A gap between theory and practice is observed for a dense version of ResNet18 trained with standard tools: the bound is 30 orders of magnitude too large on ImageNet.

Possible leads to close the gap between theory and practice. 1) Without changing the bound of Theorem 3.1, sparsity seems promising to reduce the path-norm by several orders of magnitude without changing the performance of the network. 2) Theorem 3.1 results from the worst situation (that can be met) where all the inputs activate all the paths of the network simultaneously. Bounds involving the expected path-activations could be tighter. The coordinates of $\Phi(\theta)$ are elementary bricks that can be summed to get the slopes of R_θ on the different region where R_θ is affine (Arona

et al., 2017), $\|\Phi(\theta)\|_1$ is the sum of all the bricks in absolute value, resulting in a worst-case uniform bound for all the slopes. Ideally, the bound should rather depend on the expected slopes over the different regions, weighted by the probability of falling into these regions. 3) Weight sharing may leave room for sharpened analysis (Pitas et al., 2019; Galanti et al., 2023). 4) A max pooling neuron with kernel size k only activates $1/k$ of the paths, but the bound sums the coordinates of Φ related to these k paths. This may lead to a bound k times too large in general (or even more in the presence of multiple maxpooling layers). 5) Possible bounds involving the L^q path-norm for $q > 1$ deserve a particular attention, since numerical evaluations show that they are several orders of magnitude below the L^1 norm.

Extensions to other architectures. Despite its applicability to a wide range of standard modern networks, the generalization bound in Theorem 3.1 does not cover networks with other activations than ReLU, identity, and max. The same proof technique could extend the generalization bound to new activations 1) that are positively homogeneous so that the weights can be rescaled without changing the associated function, and 2) that satisfy a contraction lemma similar to the one established here for ReLU and max neurons (typically requiring Lipschitz activation). Plausible candidates for satisfying both properties are Leaky ReLU and GroupSort (Anil et al., 2019). For SiLU (for Efficient Nets) and Hardswish (for MobileNet-V3), parts of the technical lemmas related to contraction may extend since they are Lipschitz, but rescalings would not as these activations are not positively homogeneous.

ACKNOWLEDGMENTS

This work was supported in part by the AllegroAssai ANR-19-CHIA-0009 and NuSCAP ANR-20-CE48-0014 projects of the French Agence Nationale de la Recherche.

The authors thank the Blaise Pascal Center for the computational means. It uses the SIDUS (Quemener & Corvellec, 2013) solution developed by Emmanuel Quemener.

REFERENCES

- Pierre Alquier. User-friendly introduction to PAC-Bayes bounds. *CoRR*, abs/2110.11216, 2021. URL <https://arxiv.org/abs/2110.11216>.
- Cem Anil, James Lucas, and Roger B. Grosse. Sorting out Lipschitz function approximation. In Kamalika Chaudhuri and Ruslan Salakhutdinov (eds.), *Proceedings of the 36th International Conference on Machine Learning, ICML 2019, 9-15 June 2019, Long Beach, California, USA*, volume 97 of *Proceedings of Machine Learning Research*, pp. 291–301. PMLR, 2019. URL <http://proceedings.mlr.press/v97/anil19a.html>.
- Raman Arora, Amitabh Basu, Poorya Mianjy, and Anirbit Mukherjee. Understanding deep neural networks with rectified linear units. *Electron. Colloquium Comput. Complex.*, 24:98, 2017. URL <https://eccc.weizmann.ac.il/report/2017/098>.
- Francis Bach. Learning from first principles. URL https://www.di.ens.fr/~fbach/ltfp_book.pdf.
- Francis R. Bach. Breaking the curse of dimensionality with convex neural networks. *J. Mach. Learn. Res.*, 18:19:1–19:53, 2017. URL <http://jmlr.org/papers/v18/14-546.html>.
- Andrew R. Barron and Jason M. Klusowski. Complexity, statistical risk, and metric entropy of deep nets using total path variation. *CoRR*, abs/1902.00800, 2019. URL <http://arxiv.org/abs/1902.00800>.
- Peter L. Bartlett and Shahar Mendelson. Rademacher and Gaussian complexities: Risk bounds and structural results. *J. Mach. Learn. Res.*, 3:463–482, 2002. URL <http://jmlr.org/papers/v3/bartlett02a.html>.
- Peter L. Bartlett, Dylan J. Foster, and Matus Telgarsky. Spectrally-normalized margin bounds for neural networks. In Isabelle Guyon, Ulrike von Luxburg, Samy Bengio,

- Hanna M. Wallach, Rob Fergus, S. V. N. Vishwanathan, and Roman Garnett (eds.), *Advances in Neural Information Processing Systems 30: Annual Conference on Neural Information Processing Systems 2017, December 4-9, 2017, Long Beach, CA, USA*, pp. 6240–6249, 2017. URL <https://proceedings.neurips.cc/paper/2017/hash/b22b257ad0519d4500539da3c8bcf4dd-Abstract.html>.
- Stéphane Boucheron, Gábor Lugosi, and Pascal Massart. *Concentration inequalities*. Oxford University Press, Oxford, 2013. ISBN 978-0-19-953525-5. doi: 10.1093/acprof:oso/9780199535255.001.0001. URL <https://doi-org.acces.bibliotheque-diderot.fr/10.1093/acprof:oso/9780199535255.001.0001>. A nonasymptotic theory of independence, With a foreword by Michel Ledoux.
- Gintare Karolina Dziugaite. *Revisiting Generalization for Deep Learning: PAC-Bayes, Flat Minima, and Generative Models*. PhD thesis, Department of Engineering University of Cambridge, 2018.
- Gintare Karolina Dziugaite and Daniel M. Roy. Computing nonvacuous generalization bounds for deep (stochastic) neural networks with many more parameters than training data. In Gal Elidan, Kristian Kersting, and Alexander Ihler (eds.), *Proceedings of the Thirty-Third Conference on Uncertainty in Artificial Intelligence, UAI 2017, Sydney, Australia, August 11-15, 2017*. AUAI Press, 2017. URL <http://auai.org/uai2017/proceedings/papers/173.pdf>.
- Gintare Karolina Dziugaite, Alexandre Drouin, Brady Neal, Nitarshan Rajkumar, Ethan Caballero, Linbo Wang, Ioannis Mitliagkas, and Daniel M. Roy. In search of robust measures of generalization. In Hugo Larochelle, Marc’Aurelio Ranzato, Raia Hadsell, Maria-Florina Balcan, and Hsuan-Tien Lin (eds.), *Advances in Neural Information Processing Systems 33: Annual Conference on Neural Information Processing Systems 2020, NeurIPS 2020, December 6-12, 2020, virtual*, 2020. URL <https://proceedings.neurips.cc/paper/2020/hash/86d7c8a08b4aa1bc7c599473f5ddda-Abstract.html>.
- Weinan E, Chao Ma, and Lei Wu. The Barron space and the flow-induced function spaces for neural network models. *Constr. Approx.*, 55(1):369–406, 2022. ISSN 0176-4276. doi: 10.1007/s00365-021-09549-y. URL <https://doi-org.acces.bibliotheque-diderot.fr/10.1007/s00365-021-09549-y>.
- Jonathan Frankle, David J. Schwab, and Ari S. Morcos. The early phase of neural network training. In *8th International Conference on Learning Representations, ICLR 2020, Addis Ababa, Ethiopia, April 26-30, 2020*. OpenReview.net, 2020. URL <https://openreview.net/forum?id=Hk1liRNFwS>.
- Jonathan Frankle, Gintare Karolina Dziugaite, Daniel M. Roy, and Michael Carbin. Pruning neural networks at initialization: Why are we missing the mark? In *9th International Conference on Learning Representations, ICLR 2021, Virtual Event, Austria, May 3-7, 2021*. OpenReview.net, 2021. URL <https://openreview.net/forum?id=Ig-VyQc-MLK>.
- Yasutaka Furusho. *Analysis of Regularization and Optimization for Deep Learning*. PhD thesis, Nara Institute of Science and Technology, 2020.
- Tomer Galanti, Mengjia Xu, Liane Galanti, and Tomaso A. Poggio. Norm-based generalization bounds for compositionally sparse neural networks. *CoRR*, abs/2301.12033, 2023. doi: 10.48550/arXiv.2301.12033. URL <https://doi.org/10.48550/arXiv.2301.12033>.
- Noah Golowich, Alexander Rakhlin, and Ohad Shamir. Size-independent sample complexity of neural networks. In Sébastien Bubeck, Vianney Perchet, and Philippe Rigollet (eds.), *Conference on Learning Theory, COLT 2018, Stockholm, Sweden, 6-9 July 2018*, volume 75 of *Proceedings of Machine Learning Research*, pp. 297–299. PMLR, 2018. URL <http://proceedings.mlr.press/v75/golowich18a.html>.
- Antoine Gonon, Nicolas Brisebarre, Rémi Gribonval, and Elisa Riccietti. Approximation speed of quantized vs. unquantized ReLU neural networks and beyond. *IEEE Transactions on Information Theory*, 2023. doi: 10.1109/TIT.2023.3240360.
- Benjamin Guedj. A primer on PAC-Bayesian learning. *CoRR*, abs/1901.05353, 2019. URL <http://arxiv.org/abs/1901.05353>.

- Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *2016 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2016, Las Vegas, NV, USA, June 27-30, 2016*, pp. 770–778. IEEE Computer Society, 2016. doi: 10.1109/CVPR.2016.90. URL <https://doi.org/10.1109/CVPR.2016.90>.
- Yiding Jiang, Behnam Neyshabur, Hossein Mobahi, Dilip Krishnan, and Samy Bengio. Fantastic generalization measures and where to find them. In *8th International Conference on Learning Representations, ICLR 2020, Addis Ababa, Ethiopia, April 26-30, 2020*. OpenReview.net, 2020. URL <https://openreview.net/forum?id=SJgIPJBFvH>.
- Sham M. Kakade, Karthik Sridharan, and Ambuj Tewari. On the complexity of linear prediction: Risk bounds, margin bounds, and regularization. In Daphne Koller, Dale Schuurmans, Yoshua Bengio, and Léon Bottou (eds.), *Advances in Neural Information Processing Systems 21, Proceedings of the Twenty-Second Annual Conference on Neural Information Processing Systems, Vancouver, British Columbia, Canada, December 8-11, 2008*, pp. 793–800. Curran Associates, Inc., 2008. URL <https://proceedings.neurips.cc/paper/2008/hash/5b69b9cb83065d403869739ae7f0995e-Abstract.html>.
- Kenji Kawaguchi, Leslie Pack Kaelbling, and Yoshua Bengio. Generalization in deep learning. *CoRR*, abs/1710.05468, 2017. URL <http://arxiv.org/abs/1710.05468>.
- Michel Ledoux and Michel Talagrand. *Probability in Banach spaces*, volume 23 of *Ergebnisse der Mathematik und ihrer Grenzgebiete (3) [Results in Mathematics and Related Areas (3)]*. Springer-Verlag, Berlin, 1991. ISBN 3-540-52013-9. doi: 10.1007/978-3-642-20212-4. URL <https://doi.org/10.1007/978-3-642-20212-4>. Isoperimetry and processes.
- Sibylle Marcotte, Rémi Gribonval, and Gabriel Peyré. Abide by the law and follow the flow: Conservation laws for gradient flows. *CoRR*, abs/2307.00144, 2023. doi: 10.48550/arXiv.2307.00144. URL <https://doi.org/10.48550/arXiv.2307.00144>.
- Andreas Maurer. A vector-contraction inequality for rademacher complexities. In Ronald Ortner, Hans Ulrich Simon, and Sandra Zilles (eds.), *Algorithmic Learning Theory - 27th International Conference, ALT 2016, Bari, Italy, October 19-21, 2016, Proceedings*, volume 9925 of *Lecture Notes in Computer Science*, pp. 3–17, 2016. doi: 10.1007/978-3-319-46379-7_1. URL https://doi.org/10.1007/978-3-319-46379-7_1.
- Vaishnavh Nagarajan and J. Zico Kolter. Uniform convergence may be unable to explain generalization in deep learning. In Hanna M. Wallach, Hugo Larochelle, Alina Beygelzimer, Florence d’Alché-Buc, Emily B. Fox, and Roman Garnett (eds.), *Advances in Neural Information Processing Systems 32: Annual Conference on Neural Information Processing Systems 2019, NeurIPS 2019, December 8-14, 2019, Vancouver, BC, Canada*, pp. 11611–11622, 2019. URL <https://proceedings.neurips.cc/paper/2019/hash/05e97c207235d63ceb1db43c60db7bbb-Abstract.html>.
- Behnam Neyshabur. Implicit regularization in deep learning. *CoRR*, abs/1709.01953, 2017. URL <http://arxiv.org/abs/1709.01953>.
- Behnam Neyshabur, Ryota Tomioka, and Nathan Srebro. Norm-based capacity control in neural networks. In Peter Grünwald, Elad Hazan, and Satyen Kale (eds.), *Proceedings of The 28th Conference on Learning Theory, COLT 2015, Paris, France, July 3-6, 2015*, volume 40 of *JMLR Workshop and Conference Proceedings*, pp. 1376–1401. JMLR.org, 2015. URL <http://proceedings.mlr.press/v40/Neyshabur15.html>.
- Behnam Neyshabur, Srinadh Bhojanapalli, David McAllester, and Nati Srebro. Exploring generalization in deep learning. In Isabelle Guyon, Ulrike von Luxburg, Samy Bengio, Hanna M. Wallach, Rob Fergus, S. V. N. Vishwanathan, and Roman Garnett (eds.), *Advances in Neural Information Processing Systems 2017, December 4-9, 2017, Long Beach, CA, USA*, pp. 5947–5956, 2017. URL <https://proceedings.neurips.cc/paper/2017/hash/10ce03aled01077e3e289f3e53c72813-Abstract.html>.

- Behnam Neyshabur, Srinadh Bhojanapalli, and Nathan Srebro. A PAC-Bayesian approach to spectrally-normalized margin bounds for neural networks. In *6th International Conference on Learning Representations, ICLR 2018, Vancouver, BC, Canada, April 30 - May 3, 2018, Conference Track Proceedings*. OpenReview.net, 2018. URL https://openreview.net/forum?id=Skz_WfbCZ.
- Guillermo Valle Pérez and Ard A. Louis. Generalization bounds for deep learning. *CoRR*, abs/2012.04115, 2020. URL <https://arxiv.org/abs/2012.04115>.
- Konstantinos Pitas, Andreas Loukas, Mike Davies, and Pierre Vandergheynst. Some limitations of norm based generalization bounds in deep neural networks. *CoRR*, abs/1905.09677, 2019. URL <http://arxiv.org/abs/1905.09677>.
- E. Quemener and M. Corvellec. SIDUS—the Solution for Extreme Deduplication of an Operating System. *Linux Journal*, 2013.
- Shai Shalev-Shwartz and Shai Ben-David. *Understanding Machine Learning - From Theory to Algorithms*. Cambridge University Press, 2014. ISBN 978-1-10-705713-5. URL <http://www.cambridge.org/de/academic/subjects/computer-science/pattern-recognition-and-machine-learning/understanding-machine-learning-theory-algorithms>.
- Pierre Stock and Rémi Gribonval. An embedding of ReLU networks and an analysis of their identifiability. *Constructive Approximation*, July 2022. ISSN 1432-0940. doi: 10.1007/s00365-022-09578-1. URL <https://doi.org/10.1007/s00365-022-09578-1>.
- Ulrike von Luxburg and Olivier Bousquet. Distance-based classification with Lipschitz functions. *J. Mach. Learn. Res.*, 5:669–695, 2004. URL <http://jmlr.org/papers/volume5/luxburg04b/luxburg04b.pdf>.
- Martin J. Wainwright. *High-dimensional statistics*, volume 48 of *Cambridge Series in Statistical and Probabilistic Mathematics*. Cambridge University Press, Cambridge, 2019. ISBN 978-1-108-49802-9. doi: 10.1017/9781108627771. URL <https://doi-org.acces.bibliotheque-diderot.fr/10.1017/9781108627771>. A non-asymptotic viewpoint.
- Chiyuan Zhang, Samy Bengio, Moritz Hardt, Benjamin Recht, and Oriol Vinyals. Understanding deep learning (still) requires rethinking generalization. *Commun. ACM*, 64(3):107–115, 2021. doi: 10.1145/3446776. URL <https://doi.org/10.1145/3446776>.
- Shuxin Zheng, Qi Meng, Huishuai Zhang, Wei Chen, Nenghai Yu, and Tie-Yan Liu. Capacity control of ReLU neural networks by basis-path norm. In *The Thirty-Third AAAI Conference on Artificial Intelligence, AAAI 2019, The Thirty-First Innovative Applications of Artificial Intelligence Conference, IAAI 2019, The Ninth AAAI Symposium on Educational Advances in Artificial Intelligence, EAAI 2019, Honolulu, Hawaii, USA, January 27 - February 1, 2019*, pp. 5925–5932. AAAI Press, 2019. doi: 10.1609/aaai.v33i01.33015925. URL <https://doi.org/10.1609/aaai.v33i01.33015925>.

Supplementary material

A MODEL’S BASICS

The next definition introduces the path-embedding and the path-activations associated with the general model described in Definition 2.1.

Definition A.1 (Path-embedding and path-activations). *Consider a DAG ReLU neural network architecture G as in Definition 2.1 and parameters θ associated with G . Call a path of G any sequence of neurons v_1, \dots, v_d such that $v_i \rightarrow v_{i+1}$ is an edge. This includes paths p reduced to a single $v \in N_{out}$. Denote \mathcal{P}^G the set of paths ending at an output neuron of G . For $p \in \mathcal{P}^G$,*

$$\Phi_p(\theta) = b_{v_1} \prod_{i=1}^{d-1} \theta^{v_i \rightarrow v_{i+1}},$$

where, for practical purposes, we extended θ to input neurons v by setting $b_v = 1$, and to max pooling neurons v by setting $b_v = 0$. The path-embedding $\Phi^G(\theta)$ of θ is

$$\Phi^G(\theta) = (\Phi_p(\theta))_{p \in \mathcal{P}^G}.$$

This is often denoted Φ when the graph G is clear from the context. Moreover, given a neuron v of G , we often denote $\Phi^{\rightarrow v}$ to be the path-embedding associated with the graph deduce from G by keeping only the subgraph ending at v : every neuron that cannot reach v through the edges of G is removed as well as all its incoming and outgoing edges.

Consider an input x of G . Say that a path $p = v_1 \rightarrow \dots \rightarrow v_d$ is active on input x and parameters θ , and denote $a_p(\theta, x) = 1$, if for every ReLU neuron v along p , it holds $v(\theta, x) \geq 0$, and if for every max neuron v_i along p , neuron v_{i-1} is the first in $\text{ant}(v_i)$ in lexicographic order to satisfy $v_{i-1}(\theta, x) = \max_{v \in \text{ant}(v_i)} v(\theta, x)$. Otherwise, denote $a_p(\theta, x) = 0$. Consider a new symbol b (for bias) that is not used for denoting neurons. The path-activations matrix $\mathbf{A}(\theta, x)$ is defined as the matrix in $\mathbb{R}^{\mathcal{P} \times \overline{N_m}}$ such that for any path $p \in \mathcal{P}$ and neuron $u \in \overline{N_{in}}$

$$(\mathbf{A}(\theta, x))_{p,u} = \begin{cases} a_p(\theta, x) \mathbb{1}_{p \text{ starts at } u} & \text{if } u \in N_{in}, \\ a_p(\theta, x) & \text{otherwise when } u = b. \end{cases}$$

The next lemma shows how the path-embedding and the path-activations are an equivalent way to define the model. The proof is at the end of the section.

Lemma A.1. *Consider a model as in Definition 2.1. Then for every neuron v , every input x and every parameters θ :*

$$v(\theta, x) = \left\langle \Phi^{\rightarrow v}(\theta), \mathbf{A}^{\rightarrow v}(\theta, x) \begin{pmatrix} x \\ 1 \end{pmatrix} \right\rangle.$$

Proof of Lemma A.1. For any neuron v , denote $\mathcal{P}^{\rightarrow v}$ the set of paths ending at neuron v . We want to prove that for any neuron v :

$$\begin{aligned} v(\theta, x) &= \left\langle \Phi^{\rightarrow v}(\theta), \mathbf{A}^{\rightarrow v}(\theta, x) \begin{pmatrix} x \\ 1 \end{pmatrix} \right\rangle \\ &= \sum_{p \in \mathcal{P}^{\rightarrow v}} \Phi_p(\theta) a_p(\theta, x) x_{p_0}. \end{aligned}$$

where we denote in the proof $x_u = 1$ for any u which is not an input neuron, and where p_0 denotes the first neuron of a path p . This is true by convention for input neurons v . Indeed, considering the path $p = v$, it holds $\Phi_p(\theta) = b_v = 1$, $a_p(\theta, x) = 1$ and $v(\theta, x) = x_v = x$.

Consider now v which is not an input neuron and assume that this is true for every neuron $u \in \text{ant}(v)$. If v is an identity neuron or a ReLU neuron, then

$$\begin{aligned} v(\theta, x) &= \rho_v \left(b_v + \sum_{u \in \text{ant}(v)} \theta^{u \rightarrow v} u(\theta, x) \right) \\ &= \rho_v \left(b_v + \sum_{u \in \text{ant}(v)} \theta^{u \rightarrow v} \sum_{p \in \mathcal{P}^{\rightarrow u}} \Phi_p(\theta) a_p(\theta, x) x_{p_0} \right). \end{aligned}$$

using the assumption on the antecedents of v . For a path $\tilde{p} = p \rightarrow v$ with $p \in \mathcal{P}^{\rightarrow u}$, it holds $x_{p_0} = x_{\tilde{p}_0}$, $\Phi_{\tilde{p}}(\boldsymbol{\theta}) = \boldsymbol{\theta}^{u \rightarrow v} \Phi_p(\boldsymbol{\theta})$ and $b_v = \Phi_p(\boldsymbol{\theta}) x_{p_0}$ for the path $p = v$. The latter is indeed true because $\Phi_p(\boldsymbol{\theta}) = b_v$ and $x_{p_0} = 1$ by convention since v is not an input neuron.

If v is an identity neuron, then $a_{\tilde{p}}(\boldsymbol{\theta}, x) = a_p(\boldsymbol{\theta}, x)$ and $a_v(\boldsymbol{\theta}, x) = 1$. Since

$$\mathcal{P}^{\rightarrow v} = \{v\} \cup \left(\bigcup_{u \in \text{ant}(v)} \{p \rightarrow v, p \in \mathcal{P}^{\rightarrow u}\} \right),$$

this yields the result in the case of an identity neuron v . If v is a ReLU neuron, then

$$a_{\tilde{p}}(\boldsymbol{\theta}, x) = a_p(\boldsymbol{\theta}, x) \mathbb{1}_{v(\boldsymbol{\theta}, x) \geq 0}$$

and for the path $p = v$, it holds $a_p(\boldsymbol{\theta}, x) = \mathbb{1}_{v(\boldsymbol{\theta}, x) \geq 0}$ so that once again the result holds true:

$$\begin{aligned} v(\boldsymbol{\theta}, x) &= \rho_v \left(b_v + \sum_{u \in \text{ant}(v)} \boldsymbol{\theta}^{u \rightarrow v} \sum_{p \in \mathcal{P}^{\rightarrow u}} \Phi_p(\boldsymbol{\theta}) a_p(\boldsymbol{\theta}, x) x_{p_0} \right) \\ &= \mathbb{1}_{v(\boldsymbol{\theta}, x) \geq 0} \left(b_v + \sum_{u \in \text{ant}(v)} \sum_{p \in \mathcal{P}^{\rightarrow u}} \boldsymbol{\theta}^{u \rightarrow v} \Phi_p(\boldsymbol{\theta}) a_p(\boldsymbol{\theta}, x) x_{p_0} \right) \\ &= \sum_{p \in \mathcal{P}^{\rightarrow v}} \Phi_p(\boldsymbol{\theta}) a_p(\boldsymbol{\theta}, x) x_{p_0}. \end{aligned}$$

If v is a max pooling neuron, then

$$\begin{aligned} v(\boldsymbol{\theta}, x) &= \max_{u \in \text{ant}(v)} \boldsymbol{\theta}^{u \rightarrow v} u(\boldsymbol{\theta}, x) + b_v \\ &= \max_{u \in \text{ant}(v)} \boldsymbol{\theta}^{u \rightarrow v} \sum_{p \in \mathcal{P}^{\rightarrow u}} \Phi_p(\boldsymbol{\theta}) a_p(\boldsymbol{\theta}, x) x_{p_0} + b_v \end{aligned}$$

with ties decided by lexicographic order. Since for any $\tilde{p} = p \rightarrow v$ with $\tilde{p} \in \mathcal{P}^{\rightarrow u}$, it holds

$$a_{\tilde{p}}(\boldsymbol{\theta}, x) = a_p(\boldsymbol{\theta}, x) \mathbb{1}_{u \text{ realizes the max in lexicographic order for } (\boldsymbol{\theta}, x)}$$

thus once again the claim holds for v . This proves the result. \square

A straightforward consequence of Lemma A.1 is the Lipschitz bound $\|R_{\boldsymbol{\theta}}(x) - R_{\boldsymbol{\theta}}(x')\|_1 \leq \|\Phi(\boldsymbol{\theta})\|_1 \|x - x'\|_{\infty}$. This fact is already mentioned in the case of feedforward neural networks without biases in (Neyshabur, 2017, before Section 3.4), and proven in (Furusho, 2020, Theorem 5).

Proof of the Lipschitz property. Consider parameters $\boldsymbol{\theta}$. Consider inputs x, x' with the same path-activations with respect to $\boldsymbol{\theta}$: $\mathbf{A}(\boldsymbol{\theta}, x) = \mathbf{A}(\boldsymbol{\theta}, x')$. Then:

$$\begin{aligned} \|R_{\boldsymbol{\theta}}(x) - R_{\boldsymbol{\theta}}(x')\|_1 &\stackrel{\text{Lemma A.1}}{=} \sum_{v \in N_{\text{out}}} \left| \left\langle \Phi^{\rightarrow v}(\boldsymbol{\theta}), \mathbf{A}^{\rightarrow v}(\boldsymbol{\theta}, x) \begin{pmatrix} x \\ 1 \end{pmatrix} - \mathbf{A}^{\rightarrow v}(\boldsymbol{\theta}, x') \begin{pmatrix} x' \\ 1 \end{pmatrix} \right\rangle \right| \\ &\leq \sum_{v \in N_{\text{out}}} \|\Phi^{\rightarrow v}(\boldsymbol{\theta})\|_1 \left\| \mathbf{A}^{\rightarrow v}(\boldsymbol{\theta}, x) \begin{pmatrix} x \\ 1 \end{pmatrix} - \mathbf{A}^{\rightarrow v}(\boldsymbol{\theta}, x') \begin{pmatrix} x' \\ 1 \end{pmatrix} \right\|_{\infty} \\ &\leq \sum_{v \in N_{\text{out}}} \|\Phi^{\rightarrow v}(\boldsymbol{\theta})\|_1 \left\| \mathbf{A}^{\rightarrow v}(\boldsymbol{\theta}, x) \left(\begin{pmatrix} x \\ 1 \end{pmatrix} - \begin{pmatrix} x' \\ 1 \end{pmatrix} \right) \right\|_{\infty} \\ &\leq \sum_{v \in N_{\text{out}}} \|\Phi^{\rightarrow v}(\boldsymbol{\theta})\|_1 \|x - x'\|_{\infty} \\ &= \|\Phi(\boldsymbol{\theta})\|_1 \|x - x'\|_{\infty}. \end{aligned}$$

We just proved the claim locally on each region where the path-activations $\mathbf{A}(\boldsymbol{\theta}, \cdot)$ are constant. Since Lipschitzness is a local property, this yields the result. \square




Figure 1: Example of a network where one must replace the max pooling neuron to compute the path-norm with a single forward pass as in Equation (4).

Another straightforward but important consequence of Lemma A.1 is that the path-norm (the norm of the path-embedding) can be computed in a single forward pass, up to replacing max pooling neurons with linear ones. Proofs are at the end of the section.

Theorem A.1. *Consider an architecture G as in Definition 2.1. Define \tilde{G} to be the same as G except for max pooling neurons for which the activation function is replaced by the identity. Consider an exponent $q \in [1, \infty)$ and arbitrary parameters θ associated with G (or equivalently \tilde{G}). Denote $|\theta|^q$ the vector deduced from θ by applying $x \mapsto |x|^q$ coordinate-wise, and by $\mathbf{1}$ the input full of ones. Then*

$$\|\Phi(\theta)\|_q^q = \|\mathbb{R}_{|\theta|^q}^{\tilde{G}}(\mathbf{1})\|_1. \quad (4)$$

Moreover, the formula is false in general if the max pooling neurons have not been replaced with identity ones (i.e. if the forward pass is done on G rather than \tilde{G}).

Proof of Theorem A.1. In the situation of Figure 1, Equation (4) is false if the max pooling neurons have not been replaced with identity ones as the forward pass would give 1 while the path-norm is equal to 2. Let us see now Equation (4) holds when considering \tilde{G} .

Denote $\mathcal{P}^{\rightarrow v}$ the set of paths ending at a neuron v . Denote by $\Phi^{\tilde{G}}$ and $a^{\tilde{G}}$ the path-embedding and the path-activations associated with \tilde{G} . According to Lemma A.1, it holds for every output neuron of \tilde{G}

$$(R_{|\theta|^q}^{\tilde{G}}(x))_v = \sum_{p \in \mathcal{P}^{\rightarrow v}} \Phi_p^{\tilde{G}}(|\theta|^q) a_p^{\tilde{G}}(|\theta|^q, x) x_{p_0}.$$

Since \tilde{G} has only identity or ReLU neurons, and since the parameters $|\theta|^q$ are non-negative, then for every input x with non-negative coordinates, a simple induction on the neurons shows that for every neuron u , it holds

$$u(|\theta|^q, x) \geq 0.$$

Thus, every path p is active, meaning that $a_p^{\tilde{G}}(|\theta|^q, x) = 1$ for every path p and every non-negative input x . Moreover, $\Phi_p^{\tilde{G}}(|\theta|^q) = \left| b_{p_0} \prod_{u \rightarrow v \in p} \theta^{u \rightarrow v} \right|^q = |\Phi_p^G(\theta)|^q$. Denoting by \mathcal{P} the set of paths ending at an output neuron, summing the previous terms over all output neurons concludes

$$\|\mathbb{R}_{|\theta|^q}^{\tilde{G}}(\mathbf{1})\|_1 = \sum_{p \in \mathcal{P}} |\Phi_p^G(\theta)|^q = \|\Phi(\theta)\|_q^q.$$

□

B RELATION BETWEEN PATH-NORMS AND PRODUCTS OF OPERATORS' NORMS

Feedforward ReLU networks. For simple models of the form $R_\theta(x) = M_D \text{ReLU}(M_{D-1} \dots \text{ReLU}(M_1 x))$, it is known that $\|\Phi(\theta)\|_q \leq \prod_{\ell=1}^D \|M_\ell\|_{q, \infty}$ (where $\|M\|_{q, \infty}$ is the maximum L^q norm of a row of matrix M) (Neyshabur et al., 2015, Theorem 5). Theorem B.1 below generalizes this result to the case of an arbitrary DAG (that may include max pooling, average pooling, skip connections) with biases. The rescaling of θ that makes it an equality without changing R_θ is given by Algorithm 1.

Algorithm 1. Algorithm 1 rescales θ while preserving R_θ because for any neuron $u \notin N_{\text{in}} \cup N_{\text{out}}$, the activation function ρ_u is positively homogeneous: $\rho_u(\lambda x) = \lambda \rho_u(x)$ for every $\lambda > 0$. Thus, $\lambda \rho_u(\frac{1}{\lambda}x) = \rho_u(x)$. Let us also give some more remarks about this algorithm, which is used here for the case of equality, and in the proof of the generalization bound. The first line of the algorithm considers a topological sorting of the neurons, *i.e.* an order on the neurons such that if $u \rightarrow v$ is an edge then u comes before v in this ordering. Such an order always exists (and it can be computed in linear time). Moreover, note that a classical max pooling neuron v has not anymore its incoming weights equal to one after rescaling, in general. This has no incidence on the validity of the generalization bound on classical max pooling neurons: rescaling is only used in the proof to reduce to another representation of the parameters that realize the same function and that is more handy to work with.

Algorithm 1 Normalization of parameters for norm $q \in [1, \infty)$

```

1: Consider a topological sorting  $v_1, \dots, v_k$  of the neurons
2: for  $v = v_1, \dots, v_k$  do
3:   if  $v \notin N_{\text{in}} \cup N_{\text{out}}$  then
4:      $\lambda_v = (\|\theta^{\rightarrow v}\|_q^q + |b_v|^q)^{1/q}$ 
5:     if  $\lambda_v = 0$  then
6:        $\theta^{v \rightarrow} = 0$ 
7:     else
8:        $\theta^{\rightarrow v} \leftarrow \frac{1}{\lambda_v} \times \theta^{\rightarrow v}$  ▷ normalize incoming weights
9:        $b_v \leftarrow \frac{1}{\lambda_v} \times b_v$  ▷ normalize bias
10:       $\theta^{v \rightarrow} \leftarrow \lambda_v \times \theta^{v \rightarrow}$  ▷ rescale outgoing weights to preserve the function  $R_\theta$ 

```

Coming back to the comparison between the path-norm and the product of operators' norms, first, we introduce the equivalent of the product of operators' norms when neurons are not regrouped in layers. Note that for the simple feedforward model as above, it holds $\prod_{d=1}^D \|M_d\|_{q,\infty} = \max_{u_0 \rightarrow \dots \rightarrow u_D, u_0 \in N_{\text{in}}} \prod_{d=1}^D \|\theta^{\rightarrow u_d}\|_q$. Indeed, all the neurons of two consecutive layers are connected, so the product of the maximum of L^q -norms over layers is also the maximum over all paths of the product of L^q -norms.

General DAG ReLU network. Consider now the case of a general DAG ReLU network. For practical purposes, we extend the parameters θ to input neurons v by setting $b_v = 1$ and to max pooling neurons by setting $b_v = 0$. Consider for every path $u_0 \rightarrow \dots \rightarrow u_D$ the quantity

$$\Pi_q(u_0 \rightarrow \dots \rightarrow u_D) := \left(\sum_{d=0}^D |b_{u_d}|^q \prod_{k=d+1}^D \|\theta^{\rightarrow u_k}\|_q^q \right)^{1/q},$$

with the convention that an empty product is equal to one. Note that when there are no biases, it holds $\Pi_q(u_0 \rightarrow \dots \rightarrow u_D) = \prod_{d=1}^D \|\theta^{\rightarrow u_d}\|_q$, and taking the maximum of this product over all paths $u_0 \rightarrow \dots \rightarrow u_D$ recovers $\prod_{d=1}^D \|M_d\|_{q,\infty}$. The next theorem shows that the L^q path-norm is the minimum of this complexity measure over all possible rescalings of the parameters θ that leave invariant the associated function R_θ . The proof is in Appendix B.

Theorem B.1. For every parameters θ ,

$$\|\Phi(\theta)\|_q \leq \left(\sum_{v \in N_{\text{out}}} |b_v|^q + \|\theta^{\rightarrow v}\|_q^q \max_{\substack{D \geq 0, \\ v \in N_{\text{out}}, \\ u_0 \in N_{\text{in}}, \\ u_0 \rightarrow \dots \rightarrow u_D \rightarrow v}} \Pi_q(u_0 \rightarrow \dots \rightarrow u_D)^q \right)^{1/q}.$$

If θ has been normalized by Algorithm 1, then this is an equality with the maximum being equal to one so that it simply holds $\|\Phi(\theta)\|_q = \left(\sum_{v \in N_{\text{out}}} |b_v|^q + \|\theta^{\rightarrow v}\|_q^q \right)^{1/q}$.




Figure 2: A network for which the path-norm is zero while the product of operators' norms scales as M^2 .

Proof of Theorem B.1. Since $\Phi(\theta) = (\Phi^{\rightarrow v}(\theta))_{v \in N_{\text{out}}}$, it holds

$$\|\Phi(\theta)\|_q^q = \sum_{v \in N_{\text{out}}} \|\Phi^{\rightarrow v}(\theta)\|_q^q.$$

For any neuron $v \in N_{\text{out}}$, $\Phi^{\rightarrow v}(\theta) = \left(\begin{array}{c} (\theta^{u \rightarrow v} \Phi^{\rightarrow u}(\theta))_{u \in \text{ant}(v)} \\ b_v \end{array} \right)$ so that

$$\|\Phi^{\rightarrow v}(\theta)\|_q^q = |b_v|^q + \sum_{u \in \text{ant}(v)} |\theta^{u \rightarrow v}|^q \|\Phi^{\rightarrow u}(\theta)\|_q^q \leq |b_v|^q + \|\theta^{\rightarrow v}\|_q^q \max_{u \in \text{ant}(v)} \|\Phi^{\rightarrow u}(\theta)\|_q^q.$$

For every $u \in \text{ant}(v)$, u cannot be an output neuron since it has at least v as a successor. Thus Lemma B.2 gives:

$$\|\Phi^{\rightarrow u}(\theta)\|_q \leq \max_{\substack{D \geq 0, \\ u_0 \in N_{\text{in}}, \\ u_0 \rightarrow \dots \rightarrow u_D \rightarrow u}} \Pi_q(u_0 \rightarrow \dots \rightarrow u_D \rightarrow u).$$

Putting everything together shows the upper-bound:

$$\|\Phi(\theta)\|_q^q \leq \left(\sum_{v \in N_{\text{out}}} |b_v|^q \right) + \left(\sum_{v \in N_{\text{out}}} \|\theta^{\rightarrow v}\|_q^q \right) \max_{\substack{D \geq 0, \\ v \in N_{\text{out}}, \\ u_0 \in N_{\text{in}}, \\ u_0 \rightarrow \dots \rightarrow u_D \rightarrow v}} (\Pi_q(u_0 \rightarrow \dots \rightarrow u_D))^q.$$

We now prove the case of equality. Recall that

$$\begin{aligned} \|\Phi(\theta)\|_q^q &= \sum_{v \in N_{\text{out}}} \|\Phi^{\rightarrow v}(\theta)\|_q^q \\ &= \sum_{v \in N_{\text{out}}} |b_v|^q + \sum_{u \in \text{ant}(v)} |\theta^{u \rightarrow v}|^q \|\Phi^{\rightarrow u}(\theta)\|_q^q \end{aligned}$$

It would then be sufficient to prove that, as soon as the parameters θ have been rescaled with Algorithm 1, then $|\theta^{u \rightarrow v}| \|\Phi^{\rightarrow u}(\theta)\|_q = |\theta^{u \rightarrow v}|$ for every $v \in N_{\text{out}}$ and $u \in \text{ant}(v)$. Indeed, we would then deduce the claim by writing:

$$\begin{aligned} \|\Phi(\theta)\|_q^q &= \sum_{v \in N_{\text{out}}} |b_v|^q + \sum_{u \in \text{ant}(v)} |\theta^{u \rightarrow v}|^q \\ &= \sum_{v \in N_{\text{out}}} |b_v|^q + \|\theta^{\rightarrow v}\|_q^q. \end{aligned}$$

It now remains to see that $|\theta^{u \rightarrow v}| \|\Phi^{\rightarrow u}(\theta)\|_q = |\theta^{u \rightarrow v}|$ is a direct consequence of the next lemma.

Lemma B.1. Consider $u \in N$ and parameters θ rescaled by Algorithm 1. If $\|\Phi^{\rightarrow u}(\theta)\|_q = 0$ then $\theta^{u \rightarrow} = 0$. Otherwise, $\|\Phi^{\rightarrow u}(\theta)\|_q = 1$.

Proof of Lemma B.1. The proof is by induction on the neurons. Consider $u \in N_{\text{in}}$. Then by convention $\Phi^{\rightarrow u}(\theta) = 1$ so the claim holds true.

Consider now $u \notin N_{\text{in}}$ and assume the claim to be true for every antecedent of u . It holds:

$$\|\Phi^{\rightarrow u}(\boldsymbol{\theta})\|_q^q = |b_u|^q + \sum_{w \in \text{ant}(u)} |\boldsymbol{\theta}^{w \rightarrow u}|^q \|\Phi^{\rightarrow w}(\boldsymbol{\theta})\|_q^q.$$

A consequence of the induction hypothesis is that $|\boldsymbol{\theta}^{w \rightarrow u}|^q \|\Phi^{\rightarrow w}(\boldsymbol{\theta})\|_q^q = |\boldsymbol{\theta}^{w \rightarrow u}|^q$ for every $w \in \text{ant}(u)$. Thus

$$\begin{aligned} \|\Phi^{\rightarrow u}(\boldsymbol{\theta})\|_q^q &= |b_u|^q + \sum_{w \in \text{ant}(u)} |\boldsymbol{\theta}^{w \rightarrow u}|^q \\ &= |b_u|^q + \|\boldsymbol{\theta}^{\rightarrow u}\|_q^q. \end{aligned}$$

The latter is either equal to 0 or 1 by Lemma B.3. Moreover, when it is equal to 0, this means that in Algorithm 1, $\lambda_u = 0$ and $\boldsymbol{\theta}^{u \rightarrow} = 0$ after rescaling since it is set to zero line 6 of Algorithm 1 when u is encountered, and the coordinates of $\boldsymbol{\theta}^{u \rightarrow}$ can only be multiplied by scalars in the remaining of the algorithms so that this property stays true for the remaining of the algorithm. This proves the claim for u , and thus the induction. \square

\square

Lemma B.2. Consider an exponent $q \in [1, \infty)$. For every neuron v , it holds

$$\min_{\substack{D \geq 0, \\ u_0 \in N_{\text{in}}, \\ u_0 \rightarrow \dots \rightarrow u_D \rightarrow v}} \Pi_q(u_0 \rightarrow \dots \rightarrow u_D \rightarrow v) \leq \|\Phi^{\rightarrow v}(\boldsymbol{\theta})\|_q \leq \max_{\substack{D \geq 0, \\ u_0 \in N_{\text{in}}, \\ u_0 \rightarrow \dots \rightarrow u_D \rightarrow v}} \Pi_q(u_0 \rightarrow \dots \rightarrow u_D \rightarrow v),$$

where by convention, an empty minimum (resp. maximum) is $-\infty$ (resp. ∞) and where we define by convention $\Phi^{\rightarrow u}(\boldsymbol{\theta}) = \boldsymbol{\theta}^{\rightarrow v} = 1$ for an input neuron $u \in N_{\text{in}}$.

Proof of Lemma B.2. The proof goes by induction on a topological sorting of the graph. The first neurons of the sorting are the neurons without antecedents, *i.e.* the input neurons by definition. The inequality is true for such neurons since it writes $-\infty \leq 1 \leq \infty$ by convention. Indeed $\|\Phi^{\rightarrow v}(\boldsymbol{\theta})\|_q = 1$, and the minimum and maximum are empty.

Consider a neuron $v \notin N_{\text{in}}$ and assume that this is true for every neuron before v in a topological sorting of the graph. By definition,

$$\Phi^{\rightarrow v}(\boldsymbol{\theta}) = \begin{pmatrix} (\boldsymbol{\theta}^{u \rightarrow v} \Phi^{\rightarrow u}(\boldsymbol{\theta}))_{u \in \text{ant}(v)} \\ b_v \end{pmatrix}$$

so that

$$\|\Phi^{\rightarrow v}(\boldsymbol{\theta})\|_q^q = |b_v|^q + \sum_{u \in \text{ant}(v)} |\boldsymbol{\theta}^{u \rightarrow v}|^q \|\Phi^{\rightarrow u}(\boldsymbol{\theta})\|_q^q.$$

Thus

$$|b_v|^q + \|\boldsymbol{\theta}^{\rightarrow v}\|_q^q \min_{u \in \text{ant}(v)} \|\Phi^{\rightarrow u}(\boldsymbol{\theta})\|_q^q \leq \|\Phi^{\rightarrow v}(\boldsymbol{\theta})\|_q^q \leq |b_v|^q + \|\boldsymbol{\theta}^{\rightarrow v}\|_q^q \max_{u \in \text{ant}(v)} \|\Phi^{\rightarrow u}(\boldsymbol{\theta})\|_q^q.$$

Every $u \in \text{ant}(v)$ must arrive before v in the topological sorting so the induction hypothesis applies to them. Thus:

$$\begin{aligned} \|\Phi^{\rightarrow v}(\boldsymbol{\theta})\|_q^q &\leq |b_v|^q + \|\boldsymbol{\theta}^{\rightarrow v}\|_q^q \max_{u \in \text{ant}(v)} \max_{\substack{D \geq 0, \\ u_0 \in N_{\text{in}}, \\ u_0 \rightarrow \dots \rightarrow u_D \rightarrow u}} (\Pi_q(u_0 \rightarrow \dots \rightarrow u_D \rightarrow u))^q \\ &= |b_v|^q + \|\boldsymbol{\theta}^{\rightarrow v}\|_q^q \max_{\substack{D \geq 1, \\ u_0 \in N_{\text{in}}, \\ u_0 \rightarrow \dots \rightarrow u_D \rightarrow v}} (\Pi_q(u_0 \rightarrow \dots \rightarrow u_D))^q \end{aligned}$$

Now, note that for any path $u_0 \rightarrow \dots \rightarrow u_D \rightarrow v$, it holds

$$|b_v|^q + \|\boldsymbol{\theta}^{\rightarrow v}\|_q^q (\Pi_q(u_0 \rightarrow \dots \rightarrow u_D))^q = (\Pi_q(u_0 \rightarrow \dots \rightarrow u_D \rightarrow v))^q.$$

Indeed, denoting $u_{D+1} = v$, it holds by definition

$$\begin{aligned}
|b_v|^q + \|\boldsymbol{\theta}^{\rightarrow v}\|_q^q (\Pi_q(u_0 \rightarrow \dots \rightarrow u_D))^q &= |b_v|^q \underbrace{\left(\prod_{k=D+2}^{D+1} \|\boldsymbol{\theta}^{\rightarrow u_k}\|_q^q \right)}_{=1 \text{ (empty product)}} + \|\boldsymbol{\theta}^{\rightarrow v}\|_q^q \sum_{d=0}^D |b_{u_d}|^q \prod_{k=d+1}^L \|\boldsymbol{\theta}^{\rightarrow u_k}\|_q^q \\
&= |b_{u_{D+1}}|^q \left(\prod_{k=D+2}^{D+1} \|\boldsymbol{\theta}^{\rightarrow u_k}\|_q^q \right) + \sum_{d=0}^D |b_{u_d}|^q \prod_{k=d+1}^{D+1} \|\boldsymbol{\theta}^{\rightarrow u_k}\|_q^q \\
&= \sum_{\ell=0}^{D+1} |b_{u_\ell}|^q \prod_{k=\ell+1}^{D+1} \|\boldsymbol{\theta}^{\rightarrow u_k}\|_q^q \\
&= (\Pi_q(u_0 \rightarrow \dots \rightarrow u_D \rightarrow v))^q.
\end{aligned}$$

Thus

$$\|\Phi^{\rightarrow v}(\boldsymbol{\theta})\|_q^q \leq \max_{\substack{D \geq 1, \\ u_0 \in N_{\text{in}}, \\ u_0 \rightarrow \dots \rightarrow u_D \rightarrow v}} (\Pi_q(u_0 \rightarrow \dots \rightarrow u_D \rightarrow v))^q$$

and the maximum does not change if we consider $D \geq 0$ since any path going from an input neuron to v must be of length at least equal to one because v is not an input neuron itself. This proves that the upper bound by induction, and a similar argument applies for the lower bound. \square

Lemma B.3. Consider an exponent $q \in [1, \infty)$. Any output parameters $\boldsymbol{\theta}$ of Algorithm 1 are normalized, in the sense that for every neuron v which is not an output neuron, it holds:

$$\|\boldsymbol{\theta}^{\rightarrow v}\|_q^q + |b_v|^q = \begin{cases} 0 & \text{if } \|\boldsymbol{\theta}^{\rightarrow v}\|_q = b_v = 0, \\ 1 & \text{otherwise.} \end{cases}$$

Proof of Lemma B.3. It is clear that the claim holds true right after iteration of line 2 of Algorithm 1 corresponding to v . And since the last time the incoming weights and the bias of a neuron v are modified is when this is the turn of v in line 2, then the claim holds true. Indeed, the neurons are seen in an order given by a topological sorting, and given the lines 8, 9 and 10, the incoming weights of v can only be modified when this is the turn of v or one of its antecedents. But the antecedents of v come before v in any topological order, so they are seen before v line 2. Moreover, from line 9, it is clear that the bias of v can only be modified when this is the turn of v line 2. \square

C RELEVANT (AND APPARENTLY NEW) CONTRACTION LEMMAS

Lemma C.1. Consider finite sets I, V, W , and for each $v \in V$, consider a set $T^v \subset (\mathbb{R}^W)^I$ that contains 0. We denote $t = (t_i)_{i \in I} \in T^v$ with $t_i = (t_{i,w})_{w \in W} \in \mathbb{R}^W$. Consider functions $f_i : \mathbb{R}^W \rightarrow \mathbb{R}$ and a finite family $\boldsymbol{\varepsilon} = (\varepsilon_j)_{j \in J}$ of independent identically distributed Rademacher variables, with the index set J that will be clear from the context. Finally, consider a convex and non-decreasing function $G : \mathbb{R} \rightarrow \mathbb{R}$. Assume that at least one of the following setting holds.

Setting 1: scalar input case. $|W| = 1$ and for every $i \in I$, f_i is 1-Lipschitz with $f_i(0) = 0$.

Setting 2: max pooling case. $f_i = \max$.

Then:

$$\mathbb{E} \max_{v \in V} \sup_{t \in T^v} G \left(\sum_{i \in I} \varepsilon_{i,v} f_i(t_i) \right) \leq \mathbb{E} \max_{v \in V} \sup_{t \in T^v} G \left(\sum_{i \in I, w \in W} \varepsilon_{i,w,v} t_{i,w} \right). \quad (5)$$

The scalar input case generalizes a well-known scalar contraction inequality (Ledoux & Talagrand, 1991, Equation (4.20)) to the case where there is a maximum over $|V| > 1$ independent copies. Note that we could not find this result in the literature. The max pooling case proves something similar to a vector-valued contraction inequality (Maurer, 2016) that is known in the specific case of where G is the identity, and for arbitrary 1-Lipschitz functions f_i such that $f_i(0) = 0$ (with a different proof, and with a factor $\sqrt{2}$ on the right-hand side). Here, the vector-valued case we are interested in is $f_i = \max$ and $G = \exp$, which is covered by Lemma C.1. We could not find it stated elsewhere.

In the proof of Lemma C.1, we reduce to the more specific and simpler case where $|V| = 1$ and $|I| = 1$ that corresponds to the next lemma.

Lemma C.2. Consider a finite set W , a set T of elements $t = (t_1, t_2) \in \mathbb{R}^W \times \mathbb{R}$ and a function $f : \mathbb{R}^W \rightarrow \mathbb{R}$. Consider also a convex non-decreasing function $G : \mathbb{R} \rightarrow \mathbb{R}$ and a family of iid Rademacher variables $(\varepsilon_j)_{j \in J}$ where J will be clear from the context. Assume that we are in one of the two following situations.

Scalar input case. f is 1-Lipschitz, satisfies $f(0) = 0$ and has a scalar input ($|W| = 1$).

Max pooling case. f returns the maximum coordinate of its input.

Then, for every constant $c \in \mathbb{R}$, it holds

$$\mathbb{E} \max_{t \in T} \left(\sup G(\varepsilon_1 f(t_1) + t_2), c \right) \leq \mathbb{E} \max_{t \in T} \left(\sup G \left(\sum_w \varepsilon_{1,w} t_{1,w} + t_2 \right), c \right),$$

where $t_1 = (t_{1,w})_{w \in W}$.

The proof of Lemma C.2 is postponed. We now prove Lemma C.1

Proof of Lemma C.1. First, because of the Lipschitz assumptions on the f_i 's and the convexity of G , everything is measurable and the expectations are well defined.

We prove the result by reducing to the simpler case of Lemma C.2.

Reduce to the case $|V| = 1$ by conditioning and iteration.

For $v \in V$, denote $\varepsilon_{\cdot, v} := (\varepsilon_{i,v})_{i \in I}$ and $\varepsilon_{\cdot, \cdot, v} := (\varepsilon_{i,w,v})_{i \in I, w \in W}$. Define

$$h_v(\varepsilon_{\cdot, v}) := \sup_{t \in T^v} G \left(\sum_{i \in I} \varepsilon_{i,v} f_i(t_i) \right),$$

$$H_v(\varepsilon_{\cdot, \cdot, v}) := \sup_{t \in T^v} G \left(\sum_{i \in I, w \in W} \varepsilon_{i,w,v} t_{i,w} \right).$$

Equation (5) is then equivalent to

$$\mathbb{E} \max_v h_v(\varepsilon_{\cdot, v}) \leq \mathbb{E} \max_v H_v(\varepsilon_{\cdot, \cdot, v}). \quad (6)$$

To prove that, it is enough to prove that if

$$\mathbb{E} \max_v h_v(\varepsilon_{\cdot, v}) \leq \mathbb{E} \max \left(\max_{v \in V_1} h_v(\varepsilon_{\cdot, v}), \max_{v \in V_2} H_v(\varepsilon_{\cdot, \cdot, v}) \right)$$

for some partition V_1, V_2 of V , with V_2 possibly empty for the initialization of the induction, then for every $u \in V_1$:

$$\mathbb{E} \max_v h_v(\varepsilon_{\cdot, v}) \leq \mathbb{E} \max \left(\max_{v \in V_1 \setminus \{u\}} h_v(\varepsilon_{\cdot, v}), \max_{v \in V_2 \cup \{u\}} H_v(\varepsilon_{\cdot, \cdot, v}) \right).$$

Indeed, Equation (6) would come directly by induction on the size of V_2 . Consider an arbitrary partition V_1, V_2 of V , with V_2 possibly empty, and consider $u \in V_1$. It is then enough to prove that

$$\mathbb{E} \max \left(\max_{v \in V_1} h_v(\varepsilon_{\cdot, v}), \max_{v \in V_2} H_v(\varepsilon_{\cdot, \cdot, v}) \right) \leq \mathbb{E} \max \left(\max_{v \in V_1 \setminus \{u\}} h_v(\varepsilon_{\cdot, v}), \max_{v \in V_2 \cup \{u\}} H_v(\varepsilon_{\cdot, \cdot, v}) \right).$$

Define the random variable $C = C(\varepsilon) = \max(\max_{v \in V_1 \setminus \{u\}} h_v(\varepsilon_{\cdot, v}), \max_{v \in V_2} H_v(\varepsilon_{\cdot, \cdot, v}))$. Then:

$$\max \left(\max_{v \in V_1} h_v(\varepsilon_{\cdot, v}), \max_{v \in V_2} H_v(\varepsilon_{\cdot, \cdot, v}) \right) = \max(C, h_u(\varepsilon_{\cdot, u}))$$

and

$$\max \left(\max_{v \in V_1 \setminus \{u\}} h_v(\varepsilon_{\cdot, v}), \max_{v \in V_2 \cup \{u\}} H_v(\varepsilon_{\cdot, \cdot, v}) \right) = \max(C, H_u(\varepsilon_{\cdot, \cdot, u}))$$

Denote $F_c(x) = \max(c, x)$. We've just seen that it is enough to prove

$$\mathbb{E}F_{C(\boldsymbol{\varepsilon})}(h_u(\boldsymbol{\varepsilon}, u)) \leq \mathbb{E}F_{C(\boldsymbol{\varepsilon})}(H_u(\boldsymbol{\varepsilon}, u)) \quad (7)$$

where we highlight that $C(\boldsymbol{\varepsilon})$ depends on $\boldsymbol{\varepsilon}$. For a constant $c \in \mathbb{R}$, denote $e(c) = \mathbb{E}F_c(h_u(\boldsymbol{\varepsilon}, u))$ and $E(c) = \mathbb{E}F_c(H_u(\boldsymbol{\varepsilon}, u))$. Denote E_X the expectation conditioned to everything except X . Since C is independent of $\boldsymbol{\varepsilon}, u$ and $\boldsymbol{\varepsilon}, \cdot, u$, it holds

$$\mathbb{E}_{\boldsymbol{\varepsilon}, u} F_{C(\boldsymbol{\varepsilon})}(h_u(\boldsymbol{\varepsilon}, u)) = e(C(\boldsymbol{\varepsilon}))$$

and

$$\mathbb{E}_{\boldsymbol{\varepsilon}, \cdot, u} F_{C(\boldsymbol{\varepsilon})}(H_u(\boldsymbol{\varepsilon}, u)) = E(C(\boldsymbol{\varepsilon})).$$

By the law of total expectation:

$$\mathbb{E}F_C(h_u(\boldsymbol{\varepsilon}, u)) = \mathbb{E}\mathbb{E}_{\boldsymbol{\varepsilon}, u} F_C(h_u(\boldsymbol{\varepsilon}, u)) = \mathbb{E}e(C(\boldsymbol{\varepsilon}))$$

and

$$\mathbb{E}F_C(H_u(\boldsymbol{\varepsilon}, u)) = \mathbb{E}\mathbb{E}_{\boldsymbol{\varepsilon}, \cdot, u} F_C(H_u(\boldsymbol{\varepsilon}, u)) = \mathbb{E}E(C(\boldsymbol{\varepsilon})).$$

To prove Equation (7), it is then enough to prove that $e(c) \leq E(c)$ for each $c \in \mathbb{R}$, that is:

$$\mathbb{E}F_c \left(\sup_{t \in T^u} G \left(\sum_{i \in I} \varepsilon_{i,u} f_i(t_i) \right) \right) \leq \mathbb{E}F_c \left(\sup_{t \in T^u} G \left(\sum_{i \in I, w \in W} \varepsilon_{i,w,u} t_{i,w} \right) \right). \quad (8)$$

We just reduced to the case where there is a single u to consider, since all the other indices $v \in V$ have been absorbed in a constant $c \in \mathbb{R}$. In order to apply Lemma C.2, it remains to reduce to the case $|I| = 1$.

Reduce to the case $|I| = 1$ by conditioning and iteration. For any $J \subset I$, define

$$h_J(\boldsymbol{\varepsilon}, u) := \sum_{j \in J} \varepsilon_{j,u} f_j(t_j),$$

$$H_J(\boldsymbol{\varepsilon}, \cdot, u) := \sum_{j \in J, w \in W} \varepsilon_{j,w,u} t_{j,w},$$

with the convention that an empty sum is zero. Equation (8) is then equivalent to

$$\mathbb{E}F_c \left(\sup_{t \in T^u} G(h_I(\boldsymbol{\varepsilon}, u)) \right) \leq \mathbb{E}F_c \left(\sup_{t \in T^u} G(H_I(\boldsymbol{\varepsilon}, \cdot, u)) \right)$$

To prove the latter, it is then enough to prove that for every partition I_1, I_2 of I , with I_2 possibly empty, if

$$\mathbb{E}F_c \left(\sup_{t \in T^u} G(h_I(\boldsymbol{\varepsilon}, u)) \right) \leq \mathbb{E}F_c \left(\sup_{t \in T^u} G(h_{I_1}(\boldsymbol{\varepsilon}, u) + H_{I_2}(\boldsymbol{\varepsilon}, \cdot, u)) \right),$$

then for every $j \in I_1$,

$$\mathbb{E}F_c \left(\sup_{t \in T^u} G(h_I(\boldsymbol{\varepsilon}, u)) \right) \leq \mathbb{E}F_c \left(\sup_{t \in T^u} G(h_{I_1 \setminus \{j\}}(\boldsymbol{\varepsilon}, u) + H_{I_2 \cup \{j\}}(\boldsymbol{\varepsilon}, \cdot, u)) \right).$$

Indeed, Equation (8) would then come by induction on the size of I_2 . Fix an arbitrary partition I_1, I_2 of I with I_2 possibly empty, and $j \in I_1$. It is then enough to prove that

$$\begin{aligned} \mathbb{E}F_c \left(\sup_{t \in T^u} G(h_{I_1}(\boldsymbol{\varepsilon}, u) + H_{I_2}(\boldsymbol{\varepsilon}, \cdot, u)) \right) \\ \leq \mathbb{E}F_c \left(\sup_{t \in T^u} G(h_{I_1 \setminus \{j\}}(\boldsymbol{\varepsilon}, u) + H_{I_2 \cup \{j\}}(\boldsymbol{\varepsilon}, \cdot, u)) \right). \quad (9) \end{aligned}$$

For $t \in T^u$, define the random variable $\varphi(t, \boldsymbol{\varepsilon}) := h_{I_1 \setminus \{j\}}(\boldsymbol{\varepsilon}, u) + H_{I_2}(\boldsymbol{\varepsilon}, \cdot, u)$. Equation (9) is then equivalent to prove

$$\mathbb{E}F_c \left(\sup_{t \in T^u} G(h_{\{j\}}(\boldsymbol{\varepsilon}_{j,u}) + \varphi(t, \boldsymbol{\varepsilon})) \right) \leq \mathbb{E}F_c \left(\sup_{t \in T^u} G(H_{\{j\}}(\boldsymbol{\varepsilon}_{j,\cdot,u}) + \varphi(t, \boldsymbol{\varepsilon})) \right).$$

For a subset $S \subset \mathbb{R}^W \times \mathbb{R}$, denote $s = (s_1, s_2) \in S$ with $s_1 \in \mathbb{R}^W$ and $s_2 \in \mathbb{R}$, and define $e(S) = \mathbb{E}F_c(\sup_{s \in S} G(h_{\{j\}}(\varepsilon_{j,u}) + s_2))$ and $E(S) = \mathbb{E}F_c(\sup_{s \in S} G(H_{\{j\}}(\varepsilon_{j,\cdot,u}) + s_2))$. Define also the random subset $S(\varepsilon) := \{(t_j, \varphi(t, \varepsilon)), t \in T^u\}$ and denote $s_1 = t_j \in \mathbb{R}^W$ and $s_2 = \varphi(t, \delta) \in \mathbb{R}$ for $s = (t_j, \varphi(t, \delta)) \in S(\varepsilon)$. Since $\varphi(t, \varepsilon)$ is independent of $\varepsilon_{j,u}$ and $(\varepsilon_{j,w,u})_w$, it holds

$$\mathbb{E}_{\varepsilon_{j,u}} F_c \left(\sup_{t \in T^u} G(h_{\{j\}}(\varepsilon_{j,u}) + \varphi(t, \varepsilon)) \right) = e(S(\varepsilon))$$

and

$$\mathbb{E}_{(\varepsilon_{j,w,u})_w} F_c \left(\sup_{t \in T^u} G(H_{\{j\}}(\varepsilon_{j,\cdot,u}) + \varphi(t, \varepsilon)) \right) = E(S(\varepsilon)).$$

By the law of total expectation,

$$\mathbb{E}F_c \left(\sup_{t \in T^u} G(h_{\{j\}}(\varepsilon_{j,u}) + \varphi(t, \varepsilon)) \right) = \mathbb{E}e(S(\varepsilon))$$

and

$$\mathbb{E}F_c \left(\sup_{t \in T^u} G(H_{\{j\}}(\varepsilon_{j,\cdot,u}) + \varphi(t, \varepsilon)) \right) = \mathbb{E}E(S(\varepsilon)).$$

It is then enough to prove that for every subset $S \subset \mathbb{R}^W \times \mathbb{R}$, it holds $e(S) \leq E(S)$, that is:

$$\mathbb{E} \max \left(c, \sup_{s \in S} G(\varepsilon_{j,u} f_j(s_1) + s_2) \right) \leq \mathbb{E} \max \left(c, \sup_{s \in S} G \left(\sum_w \varepsilon_{j,w,u} s_{1,w} + s_2 \right) \right).$$

We just reduced to the case $|I| = 1$ since there is only the index j to consider, the others indices $i \in I \setminus \{j\}$ being absorbed in S . The latter inequality is now a direct consequence of Lemma C.2. This proves the result. \square

Proof of Lemma C.2. Write $F(x) = \max(x, c)$, so that our goal can be rewritten as

$$\mathbb{E}F \left(\sup_{t \in T} G(\varepsilon_1 f(t_1) + t_2) \right) \leq \mathbb{E}F \left(\sup_{t \in T} G \left(\sum_w \varepsilon_{1,w} t_{1,w} + t_2 \right) \right)$$

which is again equivalent to

$$\mathbb{E} \sup_{t \in T} F \circ G(\varepsilon_1 f(t_1) + t_2) \leq \mathbb{E} \sup_{t \in T} F \circ G \left(\sum_w \varepsilon_{1,w} t_{1,w} + t_2 \right) \quad (10)$$

where the supremum over T has been inverted with F because F is continuous and non-decreasing.

Scalar input case. In this case, $|W| = 1$ i.e. the inputs t_1 are scalar. Since F and G are non-decreasing and convex, then so is $F \circ G$, and Equation (10) is a direct consequence of (Ledoux & Talagrand, 1991, Equation (4.20)).

Max pooling case. In this case, $f = \max$. Computing explicitly the expectation where the only random thing is $\varepsilon_1 \in \{-1, 1\}$, the left-hand side of Equation (10) is equal to

$$\frac{1}{2} \sup_{t \in T} F \circ G(f(t_1) + t_2) + \frac{1}{2} \sup_{s \in T} F \circ G(-f(s_1) + s_2).$$

Consider $s, t \in T$. Consider $w(t)$ such that $f(t_1) = \max_{w \in W} t_{1,w} = t_{1,w(t)}$. It then holds:

$$\begin{aligned} F \circ G(f(t_1) + t_2) &=_{\varepsilon \text{ centered}} F \circ G \left(t_{1,w(t)} + \mathbb{E} \left(\sum_{w \neq w(t)} \varepsilon_{1,w} t_{1,w} \right) + t_2 \right) \\ &\leq_{\text{Jensen}} \mathbb{E}F \circ G \left(t_{1,w(t)} + \sum_{w \neq w(t)} \varepsilon_{1,w} t_{1,w} + t_2 \right). \end{aligned}$$

Moreover, $-f(s_1) = -\max_w s_{1,w} \leq -s_{1,w(t)}$ and because $F \circ G$ is non-decreasing, we also get:

$$\begin{aligned} F \circ G(-f(s_1) + s_2) &\leq F \circ G(-s_{1,w(t)} + s_2) \\ &\leq F \circ G\left(-s_{1,w(t)} + \mathbb{E}\left(\sum_{w \neq w(t)} \varepsilon_{1,w} s_{1,w}\right) + s_2\right) \\ &\leq \mathbb{E} F \circ G\left(-s_{1,w(t)} + \sum_{w \neq w(t)} \varepsilon_{1,w} s_{1,w} + s_2\right). \end{aligned}$$

At the end, we get

$$\begin{aligned} &\frac{1}{2} F \circ G(f(t_1) + t_2) + \frac{1}{2} F \circ G(-f(s_1) + s_2) \\ &\leq \frac{1}{2} \mathbb{E} F \circ G\left(t_{1,w(t)} + \sum_{w \neq w(t)} \varepsilon_{1,w} t_{1,w} + t_2\right) \\ &\quad + \frac{1}{2} \mathbb{E} F \circ G\left(-s_{1,w(t)} + \sum_{w \neq w(t)} \varepsilon_{1,w} s_{1,w} + s_2\right) \\ &\leq \frac{1}{2} \mathbb{E} \sup_{r \in T} F \circ G\left(r_{1,w(t)} + \sum_{w \neq w(t)} \varepsilon_{1,w} r_{1,w} + r_2\right) \\ &\quad + \frac{1}{2} \mathbb{E} \sup_{r \in T} F \circ G\left(-r_{1,w(t)} + \sum_{w \neq w(t)} \varepsilon_{1,w} r_{1,w} + r_2\right) \\ &= \mathbb{E} \sup_{r \in T} F \circ G\left(\varepsilon_{1,w(t)} r_{1,w(t)} + \sum_{w \neq w(t)} \varepsilon_{1,w} r_{1,w} + r_2\right) \\ &= \mathbb{E} \sup_{r \in T} F \circ G\left(\sum_w \varepsilon_{1,w} r_{1,w} + r_2\right). \end{aligned}$$

The latter is independent of s, t . Taking the supremum over all $s, t \in T$ yields Equation (10) and thus the claim. \square

D PEELING ARGUMENT

First, we state a simple lemma that will be used several times.

Lemma D.1. *Consider a vector $\varepsilon \in \mathbb{R}^n$ with iid Rademacher coordinates, meaning that $\mathbb{P}(\varepsilon_i = 1) = \mathbb{P}(\varepsilon_i = -1) = 1/2$. Consider a measurable function $G : \mathbb{R} \rightarrow \mathbb{R}_{\geq 0}$. Consider a set $X \subset \mathbb{R}^n$. Then*

$$\mathbb{E}_\varepsilon \sup_{x \in X} G\left(\left|\sum_{i=1}^n \varepsilon_i x_i\right|\right) \leq 2 \mathbb{E}_\varepsilon \sup_{x \in X} G\left(\sum_{i=1}^n \varepsilon_i x_i\right).$$

Proof of Lemma D.1. Since $G \geq 0$, it holds $G(|x|) \leq G(x) + G(-x)$. Thus

$$\mathbb{E}_\varepsilon \sup_{x \in X} G\left(\left|\sum_{i=1}^n \varepsilon_i x_i\right|\right) \leq \mathbb{E}_\varepsilon \sup_{x \in X} G\left(\sum_{i=1}^n \varepsilon_i x_i\right) + \mathbb{E}_\varepsilon \sup_{x \in X} G\left(\sum_{i=1}^n (-\varepsilon_i) x_i\right).$$

Since ε is symmetric, that is $-\varepsilon$ has the same distribution as ε , then the latter is just $2 \mathbb{E}_\varepsilon \sup_{x \in X} G\left(\sum_{i=1}^n \varepsilon_i x_i\right)$. This proves the claim. \square

Notations We now fix for all the next results of this section n vectors $x_1, \dots, x_n \in \mathbb{R}^{d_{\text{in}}}$, for some $d_{\text{in}} \in \mathbb{N}_{>0}$. We denote $x_{i,u}$ the coordinate u of x_i .

For any neural network architecture, recall that $v(\boldsymbol{\theta}, x)$ is the output of neuron v for parameters $\boldsymbol{\theta}$ and input x , and $\text{ant}^d(v)$ is the set of neurons u for which there exists a path from u to v of distance d . For a set of neurons V , denote $R_V(\boldsymbol{\theta}, x) = (v(\boldsymbol{\theta}, x))_{v \in V}$.

Introduction to peeling This section shows that some expected sum over output neurons v can be reduced to an expected maximum over $\text{ant}(v)$, and iteratively over an expected maximum over $\text{ant}^d(v)$ for increasing d 's. Eventually, the maximum is only over input neurons as soon as d is large enough. We start with the next lemma which is the initialization of the induction over d : it peels off the output neurons v to reduce to their antecedents $\text{ant}(v)$.

Lemma D.2. *Consider a neural network architecture as in Definition 2.1 with an associated set Θ of parameters $\boldsymbol{\theta}$, rescaled with Algorithm 1, and such that $\|\Phi(\boldsymbol{\theta})\|_1 \leq r$. Consider a family of independent Rademacher variables $(\varepsilon_j)_{j \in J}$ with J that will be clear from the context. Consider a non-decreasing function $G : \mathbb{R} \rightarrow \mathbb{R}$. Consider a new neuron b (for bias) and set by convention $x_b = 1$ for every input x . Then*

$$\begin{aligned} \mathbb{E}_\varepsilon G \left(\sup_{\boldsymbol{\theta} \in \Theta} \sum_{\substack{i=1, \dots, n, \\ v \in N_{\text{out}}}} \varepsilon_{i,v} v(\boldsymbol{\theta}, x_i) \right) \\ \leq \mathbb{E}_\varepsilon G \left(r \max_{v \in N_{\text{out}}} \max_{u \in (\text{ant}(v) \cap N_{\text{in}}) \cup \{b\}} \left| \sum_{i=1}^n \varepsilon_{i,v} x_{i,u} \right| \right) \\ + \mathbb{E}_\varepsilon G \left(r \max_{v \in N_{\text{out}}} \max_{u \in \text{ant}(v) \setminus N_{\text{in}}} \sup_{\boldsymbol{\theta}} \left| \sum_{i=1}^n \varepsilon_{i,v} u(\boldsymbol{\theta}, x_i) \right| \right) \end{aligned}$$

Proof of Lemma D.2. Recall that for a set of neurons V , we denote $R_V(\boldsymbol{\theta}, x) = (v(\boldsymbol{\theta}, x))_{v \in V}$. Recall that by Definition 2.1, output neurons are identity neurons so that for every $v \in N_{\text{out}}$, $\boldsymbol{\theta} \in \Theta$ and every input x :

$$v(\boldsymbol{\theta}, x) = \left\langle \left(\begin{array}{c} \boldsymbol{\theta}^{\rightarrow v} \\ b_v \end{array} \right), \left(\begin{array}{c} R_{\text{ant}(v)}(\boldsymbol{\theta}, x) \\ 1 \end{array} \right) \right\rangle.$$

We overload the symbol b to make it represent a new neuron that computes the constant function equal to one: $b(\boldsymbol{\theta}, x) = 1$.

$$\begin{aligned} \mathbb{E}_\varepsilon G \left(\sup_{\boldsymbol{\theta}} \sum_{\substack{i=1, \dots, n \\ v \in N_{\text{out}}}} \varepsilon_{i,v} v(\boldsymbol{\theta}, x_i) \right) \\ = \mathbb{E}_\varepsilon G \left(\sup_{\boldsymbol{\theta}} \sum_{v \in N_{\text{out}}} \left(\left\langle \left(\begin{array}{c} \boldsymbol{\theta}^{\rightarrow v} \\ b_v \end{array} \right), \sum_{i=1}^n \varepsilon_{i,v} \left(\begin{array}{c} R_{\text{ant}(v)}(\boldsymbol{\theta}, x_i) \\ 1 \end{array} \right) \right\rangle \right) \right) \\ \stackrel{\text{H\"older}}{\leq} \mathbb{E}_\varepsilon G \left(\sup_{\boldsymbol{\theta}} \left(\underbrace{\left(\sum_{v \in N_{\text{out}}} \|\boldsymbol{\theta}^{\rightarrow v}\|_1 + |b_v|}_{=\|\Phi(\boldsymbol{\theta})\|_1 \leq r \text{ (rescaled, Theorem B.1)}} \right) \max_{v \in N_{\text{out}}} \left(\left| \sum_{i=1}^n \varepsilon_{i,v} \right|, \max_{u \in \text{ant}(v)} \left| \sum_{i=1}^n \varepsilon_{i,v} u(\boldsymbol{\theta}, x_i) \right| \right) \right) \right) \\ \leq \mathbb{E}_\varepsilon G \left(r \max_{v \in N_{\text{out}}} \max_{u \in \text{ant}(v) \cup \{b\}} \sup_{\boldsymbol{\theta}} \left| \sum_{i=1}^n \varepsilon_{i,v} u(\boldsymbol{\theta}, x_i) \right| \right). \end{aligned}$$

Everything is non-negative so the maximum over $u \in \text{ant}(v) \cup \{b\}$ is smaller than the sum of the maxima over $u \in (\text{ant}(v) \cap N_{\text{in}}) \cup \{b\}$ and $u \in \text{ant}(v) \setminus N_{\text{in}}$. Note that when u is an input neuron, it simply holds $u(\boldsymbol{\theta}, x_i) = x_{i,u}$. This proves the result. \square

We now show how to peel neurons to reduce the maximum over $\text{ant}^d(v)$ to $\text{ant}^{d+1}(v)$. Later, we will repeat that until the maximum is only on input neurons.

Lemma D.3. Consider a neural network architecture with an associated set Θ of parameters θ rescaled by Algorithm 1. Consider a family of independent Rademacher variables $(\varepsilon_j)_{j \in J}$ with J that will be clear from the context. Consider arbitrary $M, d \in \mathbb{N}$ and a convex non-decreasing function $G : \mathbb{R} \rightarrow \mathbb{R}_{\geq 0}$. Take a symbol b (for bias) which does not correspond to a neuron ($b \notin N$) and set by convention $x_b = 1$ for every input x . Denote k the maximal kernel size of the network (i.e. the maximum of $|\text{ant}(u)|$ over every max pooling neuron u). Then:

$$\begin{aligned} \mathbb{E}_\varepsilon G \left(\max_{m=1, \dots, M} \max_{v \in N_{\text{out}}, u \in \text{ant}^d(v) \setminus N_{\text{in}}} \sup_{\theta} \left| \sum_{i=1}^n \varepsilon_{i,v,m} u(\theta, x_i) \right| \right) \\ \leq 5 \mathbb{E}_\varepsilon G \left(\max_{m=1, \dots, kM} \max_{v \in N_{\text{out}}, u \in (\text{ant}^{d+1}(v) \cap N_m) \cup \{b\}} \left| \sum_{i=1}^n \varepsilon_{i,v,m} x_{i,u} \right| \right) \\ + 5 \mathbb{E}_\varepsilon G \left(\max_{m=1, \dots, kM} \max_{v \in N_{\text{out}}, u \in \text{ant}^{d+1}(v) \setminus N_{\text{in}}} \sup_{\theta} \left| \sum_{i=1}^n \varepsilon_{i,v,m} u(\theta, x_i) \right| \right) \end{aligned}$$

Proof. Step 1: split the neurons depending on their activation function. In the term that we want to bound from above, the neurons u compute something of the form $\rho_u(\dots)$ where ρ_u is the activation associated with u which is 1-Lipschitz and satisfy $\rho_u(0) = 0$. The first step of the proof is to get rid of ρ_u using a contraction lemma of the type (Ledoux & Talagrand, 1991, Theorem 4.12). However, here, the function ρ_u depends on the neuron u , what we are taking a maximum over so that classical contraction lemmas do not apply directly. To resolve this first obstacle, we split the neurons according to their activation function. Below, we highlight in blue what is important and/or the changes from one line to another. Denote N_ρ the neurons that have ρ as their associated activation function, and the term with a maximum over all $u \in N_\rho$ is denoted:

$$e(\rho) := \mathbb{E}_\varepsilon G \left(\max_{m=1, \dots, M} \max_{v \in N_{\text{out}}, u \in (\text{ant}^d(v) \cap N_\rho) \setminus N_{\text{in}}} \sup_{\theta} \left| \sum_{i=1}^n \varepsilon_{i,v,m} u(\theta, x_i) \right| \right),$$

with the convention $e(\rho) = 0$ if N_ρ is empty. This yields a first bound

$$\mathbb{E}_\varepsilon G \left(\max_{m=1, \dots, M} \max_{v \in N_{\text{out}}, u \in \text{ant}^d(v) \setminus N_{\text{in}}} \sup_{\theta} \left| \sum_{i=1}^n \varepsilon_{i,v,m} u(\theta, x_i) \right| \right) \leq e(\text{ReLU}) + e(\text{id}) + e(\text{max}).$$

Define $E(\rho)$ to be the same thing as $e(\rho)$ but without the absolute values:

$$E(\rho) := \mathbb{E}_\varepsilon G \left(\max_{m=1, \dots, M} \max_{v \in N_{\text{out}}, u \in (\text{ant}^d(v) \cap N_\rho) \setminus N_{\text{in}}} \sup_{\theta} \sum_{i=1}^n \varepsilon_{i,v,m} u(\theta, x_i) \right).$$

Then Lemma D.1 gets rid of the absolute values by paying a factor 2:

$$e(\rho) \leq 2E(\rho).$$

We now want to bound each $E(\rho)$.

Step 2: get rid of the max and ReLU activation functions. For any max pooling neuron u that has less than k antecedents, we artificially add neurons w to $\text{ant}(u)$ to make it of cardinal k , and we set by convention $\theta^{w \rightarrow u} = 0$. We also fix an arbitrary order on the antecedents of u and wrote $\text{ant}(u)_w$ for the antecedent number w , with $R_{\text{ant}(u)_w}$ the associated function. For a ReLU or max pooling neuron u , define the pre-activation of u to be

$$\text{pre}_u(\theta, x) := \begin{cases} \left\langle \left(\begin{array}{c} \theta^{\rightarrow u} \\ b_u \end{array} \right), \left(\begin{array}{c} R_{\text{ant}(u)}(\theta, x) \\ 1 \end{array} \right) \right\rangle & \text{if } \rho_u = \text{ReLU}, \\ \left(\theta^{\text{ant}(u)_w \rightarrow u} R_{\text{ant}(u)_w}(\theta, x) \right)_{w=1, \dots, k} & \text{otherwise when } \rho_u = \text{max}. \end{cases}$$

Note that the pre-activation is defined such that $u(\boldsymbol{\theta}, x) = \rho_u(\text{pre}_u(\boldsymbol{\theta}, x))$. Consider $\rho \in \{\max, \text{ReLU}\}$. We can rewrite $E(\rho)$ in terms of the pre-activations:

$$E(\rho) = \mathbb{E}_{\boldsymbol{\varepsilon}} G \left(\max_{\substack{v \in N_{\text{out}}, \\ m=1, \dots, M}} \max_{u \in (\text{ant}^d(v) \cap N_{\rho}) \setminus N_{\text{in}}} \sup_{\boldsymbol{\theta}} \sum_{i=1}^n \varepsilon_{i,v,m} \rho(\text{pre}_u(\boldsymbol{\theta}, x_i)) \right).$$

For each $v \in N_{\text{out}}$ and $m \in \{1, \dots, M\}$, define $T^{v,m} = \{(\text{pre}_u(\boldsymbol{\theta}, x_i))_{i=1, \dots, n} \in \mathbb{R}^n, u \in (\text{ant}^d(v) \cap N_{\text{ReLU}}) \setminus N_{\text{in}}, \boldsymbol{\theta} \in \Theta\}$. We can again rewrite $E(\rho)$ as

$$E(\rho) = \mathbb{E}_{\boldsymbol{\varepsilon}} G \left(\max_{\substack{v \in N_{\text{out}}, \\ m=1, \dots, M}} \sup_{t^{v,m} \in T^{v,m}} \sum_{i=1}^n \varepsilon_{i,v,m} \rho(t_i^{v,m}) \right).$$

We now want to get rid of the activation function ρ with a contraction lemma. There is a second difficulty that prevent us from directly applying classical contraction lemmas such as (Ledoux & Talagrand, 1991, Theorem 4.12). It is the presence of a maximum over multiple copies indexed by $v \in N_{\text{out}}$ and $m \in \{1, \dots, M\}$ of a supremum that depends on iid families $(\varepsilon_{i,v,m})_{i=1 \dots n}$, since (Ledoux & Talagrand, 1991, Theorem 4.12) only deals with a single copy. This motivates the (apparently new) contraction lemma established for the occasion in Lemma C.1. Once the activation functions removed, we can conclude separately for $\rho = \text{ReLU}$, id and $\rho = \max$.

Step 3a: deal with $\rho = \max$ via rescaling. In the case $\rho = \max$, Lemma C.1 shows that

$$\begin{aligned} & \mathbb{E}_{\boldsymbol{\varepsilon}} G \left(\max_{\substack{v \in N_{\text{out}}, \\ m=1, \dots, M}} \sup_{t^{v,m} \in T^{v,m}} \sum_{i=1}^n \varepsilon_{i,v,m} \max(t_i^{v,m}) \right) \\ & \leq \mathbb{E}_{\boldsymbol{\varepsilon}} G \left(\max_{\substack{v \in N_{\text{out}}, \\ m=1, \dots, M}} \sup_{t^{v,m} \in T^{v,m}} \sum_{\substack{i=1, \dots, n, \\ w=1, \dots, k}} \varepsilon_{i,v,m,w} t_{i,w}^{v,m} \right). \end{aligned}$$

The right-hand side is equal to

$$\mathbb{E}_{\boldsymbol{\varepsilon}} G \left(\max_{\substack{v \in N_{\text{out}}, \\ m=1, \dots, M}} \sup_{u \in (\text{ant}^d(v) \cap N_{\max}) \setminus N_{\text{in}}, \boldsymbol{\theta} \in \Theta} \sum_{\substack{i=1, \dots, n, \\ w=1, \dots, k}} \varepsilon_{i,v,m,w} \boldsymbol{\theta}^{\text{ant}(u)_w \rightarrow u} R_{\text{ant}(u)_w}(\boldsymbol{\theta}, x_i) \right). \quad (11)$$

We now deal with this using the fact that the parameters are rescaled. It holds:

$$\begin{aligned} & \sum_{\substack{i=1, \dots, n, \\ w=1, \dots, k}} \varepsilon_{i,v,m,w} \boldsymbol{\theta}^{\text{ant}(u)_w \rightarrow u} R_{\text{ant}(u)_w}(\boldsymbol{\theta}, x_i) \\ & = \sum_{w=1, \dots, k} \boldsymbol{\theta}^{\text{ant}(u)_w \rightarrow u} \left(\sum_{i=1, \dots, n} \varepsilon_{i,v,m,w} R_{\text{ant}(u)_w}(\boldsymbol{\theta}, x_i) \right) \\ & \leq_{\text{H\"older}} \underbrace{\|\boldsymbol{\theta}^{\rightarrow u}\|_1}_{=0 \text{ or } 1 \text{ (no bias, rescaled Lemma B.3)}} \max_{w=1, \dots, k} \left| \sum_{i=1, \dots, n} \varepsilon_{i,v,m,w} R_{\text{ant}(u)_w}(\boldsymbol{\theta}, x_i) \right| \\ & \leq_{\text{decoupling } w \text{ and } \text{ant}(u)_w} \max_{w \in \text{ant}(u)} \max_{w'=1, \dots, k} \left| \sum_{i=1}^n \varepsilon_{i,v,m,w'} w'(\boldsymbol{\theta}, x_i) \right|. \end{aligned}$$

Thus, Equation (11) is bounded from above by

$$\mathbb{E}_{\boldsymbol{\varepsilon}} G \left(\max_{\substack{v \in N_{\text{out}}, \\ m=1, \dots, M}} \sup_{u \in (\text{ant}^d(v) \cap N_{\max}) \setminus N_{\text{in}}, \boldsymbol{\theta} \in \Theta} \max_{w \in \text{ant}(u)} \max_{w'=1, \dots, k} \left| \sum_{i=1}^n \varepsilon_{i,v,m,w'} w'(\boldsymbol{\theta}, x_i) \right| \right).$$

We can re-index the variables ε by making the third coordinate equal to the cartesian product of the current third and fourth coordinate. This absorbs the fourth coordinate in the third one, with m going from 1 to kM instead of M . Note also that for $u \in (\text{ant}^d(v) \cap N_{\max}) \setminus N_{\text{in}}$ and $w \in \text{ant}(u)$, then $w \in \text{ant}^{d+1}(v)$ so considering a maximum over $w \in \text{ant}^{d+1}(v)$ can only increase the latter expectation. Moreover, we can add a new neuron b (for bias) that computes the constant function equal to one ($b(\theta, x) = 1$) and add b to the maximum over w . At the end, Equation (11) is bounded by

$$H := \mathbb{E}_{\varepsilon} G \left(\max_{\substack{v \in N_{\text{out}}, \\ m=1, \dots, kM}} \sup_{w \in \text{ant}^{d+1}(v) \cup \{b\}} \sup_{\theta \in \Theta} \left| \sum_{i=1}^n \varepsilon_{i,v,m} w(\theta, x_i) \right| \right).$$

We now derive similar inequalities when $\rho = \text{id}$ or ReLU .

Step 3b: deal with $\rho = \text{id}, \text{ReLU}$ via rescaling. In the case $\rho = \text{ReLU}$, Lemma C.1 shows that

$$\begin{aligned} & \mathbb{E}_{\varepsilon} G \left(\max_{\substack{v \in N_{\text{out}}, \\ m=1, \dots, M}} \sup_{t^{v,m} \in T^{v,m}} \sum_{i=1}^n \varepsilon_{i,v,m} \text{ReLU}(t_i^{v,m}) \right) \\ & \leq \mathbb{E}_{\varepsilon} G \left(\max_{\substack{v \in N_{\text{out}}, \\ m=1, \dots, M}} \sup_{t^{v,m} \in T^{v,m}} \sum_{i=1, \dots, n} \varepsilon_{i,v,m} t_i^{v,m} \right). \end{aligned}$$

The difference with the max pooling case is that each $t_i^{v,m}$ is scalar so this does not introduce an additional index w to the Rademacher variables. The right-hand side can be rewritten as

$$\mathbb{E}_{\varepsilon} G \left(\max_{\substack{v \in N_{\text{out}}, \\ m=1, \dots, M}} \sup_{u \in (\text{ant}^d(v) \cap N_{\text{ReLU}}) \setminus N_{\text{in}}, \theta \in \Theta} \sum_{i=1, \dots, n} \varepsilon_{i,v,m} \left\langle \begin{pmatrix} \theta^{\rightarrow u} \\ b_u \end{pmatrix}, \begin{pmatrix} R_{\text{ant}(u)}(\theta, x_i) \\ 1 \end{pmatrix} \right\rangle \right)$$

We can only increase the latter by considering a maximum over all $u \in \text{ant}^d(v)$, not only the ones in N_{ReLU} . We also add absolute values. This is then bounded by

$$F := \mathbb{E}_{\varepsilon} G \left(\max_{\substack{v \in N_{\text{out}}, \\ m=1, \dots, M}} \sup_{u \in \text{ant}^d(v) \setminus N_{\text{in}}, \theta \in \Theta} \left| \sum_{i=1, \dots, n} \varepsilon_{i,v,m} \left\langle \begin{pmatrix} \theta^{\rightarrow u} \\ b_u \end{pmatrix}, \begin{pmatrix} R_{\text{ant}(u)}(\theta, x_i) \\ 1 \end{pmatrix} \right\rangle \right| \right). \quad (12)$$

This means that $E(\text{ReLU}) \leq F$. Let us also observe that $e(\text{id}) \leq F$. Indeed, recall that by definition

$$e(\text{id}) = \mathbb{E}_{\varepsilon} G \left(\max_{\substack{v \in N_{\text{out}}, \\ m=1, \dots, M}} \max_{u \in (\text{ant}^d(v) \cap N_{\rho}) \setminus N_{\text{in}}} \sup_{\theta} \left| \sum_{i=1}^n \varepsilon_{i,v,m} u(\theta, x_i) \right| \right).$$

We can only increase the latter expectation by considering a maximum over all $u \in \text{ant}^d(v)$. Moreover, for an identity neuron u , it holds $u(\theta, x) = \left\langle \begin{pmatrix} \theta^{\rightarrow u} \\ b_u \end{pmatrix}, \begin{pmatrix} R_{\text{ant}(u)}(\theta, x) \\ 1 \end{pmatrix} \right\rangle$. This shows that $e(\text{id}) \leq F$. It then remains to bound F using that the parameters are rescaled. Introduce a new neuron b (for bias) that computes the constant function equal to one: $b(\theta, x) = 1$. Note that

$$\begin{aligned} & \sum_{i=1, \dots, n} \varepsilon_{i,v,m} \left\langle \begin{pmatrix} \theta^{\rightarrow u} \\ b_u \end{pmatrix}, \begin{pmatrix} R_{\text{ant}(u)}(\theta, x_i) \\ 1 \end{pmatrix} \right\rangle \\ & = \left\langle \begin{pmatrix} \theta^{\rightarrow u} \\ b_u \end{pmatrix}, \sum_{i=1, \dots, n} \varepsilon_{i,v,m} \begin{pmatrix} R_{\text{ant}(u)}(\theta, x_i) \\ 1 \end{pmatrix} \right\rangle \\ & \stackrel{\text{Hölder}}{\leq} \underbrace{(\|\theta^{\rightarrow u}\|_1 + |b_u|)}_{=0 \text{ or } 1 \text{ (rescaled, Lemma B.3)}} \max_{w \in \text{ant}(u) \cup \{b\}} \left| \sum_{i=1}^n \varepsilon_{i,v,m} w(\theta, x_i) \right|. \end{aligned}$$

This shows that

$$F \leq \mathbb{E}_{\varepsilon} G \left(\max_{\substack{v \in N_{\text{out}}, \\ m=1, \dots, M}} \sup_{u \in \text{ant}^d(v) \setminus N_{\text{in}}, \theta \in \Theta} \max_{w \in \text{ant}(u) \cup \{b\}} \left| \sum_{i=1}^n \varepsilon_{i,v,m} w(\theta, x_i) \right| \right).$$

Obviously, one can introduce additional copies of ε to make the third index going from $m = 1$ to kM , and this can only increase the latter. Moreover, if $u \in \text{ant}^d(v) \setminus N_{\text{in}}$ and $w \in \text{ant}(u)$ then $w \in \text{ant}^{d+1}(v)$ so that we can consider a maximum over $w \in \text{ant}^{d+1}(v)$ and this could only increase the latter. This gives the final bound

$$\begin{aligned} F &\leq \mathbb{E}_\varepsilon G \left(\max_{\substack{v \in N_{\text{out}}, \\ m=1, \dots, kM}} \max_{w \in \text{ant}^{d+1}(v) \cup \{b\}} \sup_{\theta \in \Theta} \left| \sum_{i=1}^n \varepsilon_{i,v,m} w(\theta, x_i) \right| \right) \\ &= H. \end{aligned}$$

Step 4: putting everything together. At the end, we got the final bound

$$\begin{aligned} &\mathbb{E}_\varepsilon G \left(\max_{\substack{v \in N_{\text{out}}, \\ m=1, \dots, M}} \max_{u \in \text{ant}^d(v) \setminus N_{\text{in}}} \sup_{\theta} \left| \sum_{i=1}^n \varepsilon_{i,v,m} u(\theta, x_i) \right| \right) \\ &\leq e(\text{id}) + e(\text{ReLU}) + e(\text{max}) \\ &\leq e(\text{id}) + 2E(\text{ReLU}) + 2E(\text{max}) \\ &\leq F + 2F + 2E(\text{max}) \\ &\leq H + 2H + 2H. \end{aligned}$$

The term $5H$ can again be bounded by splitting the maximum over $w \in \text{ant}^{d+1}(v) \cup \{b\}$ between the w 's that are input neurons, and those that are not, since everything is non-negative. This yields the claim. \square

Remark D.1 (Improved dependencies on the kernel size). *Note that in the proof of Lemma D.3, the multiplication of M by k can be avoided if there are no max pooling neurons in $\text{ant}^d(v)$. Because of skip connections, even if there is a single max pooling neuron in the architecture, it can be in $\text{ant}^d(v)$ for many d 's. A more advanced version of the argument is to peel only the ReLU and identity neurons, by leaving the max pooling neurons as they are, until we reach a set of max pooling neurons large enough that we decide to peel simultaneously. This would prevent the multiplication by k every time d is increased.*

We can now state the main peeling theorem, which directly result from Lemma D.2 and Lemma D.3 by induction on d .

Theorem D.1. *Consider a neural network architecture as in Definition 2.1. Denote k its maximal kernel size (i.e. the maximum of $|\text{ant}(u)|$ over all max pooling neurons u), with $k = 1$ by convention if there is no max pooling neuron, and denote D the depth (the length of the longest path from an input to an output). Add an artificial neuron b in the input neurons N_{in} and define $x_b = 1$ for any input x . For any set of parameters Θ associated with the network, such that $\|\Phi(\theta)\|_1 \leq r$ for every $\theta \in \Theta$, it holds for every convex non-decreasing function $G : \mathbb{R} \rightarrow \mathbb{R}_{\geq 0}$*

$$\begin{aligned} &\mathbb{E}_\varepsilon G \left(\sup_{\theta \in \Theta} \sum_{\substack{i=1, \dots, n, \\ v \in N_{\text{out}}}} \varepsilon_{i,v} v(\theta, x_i) \right) \\ &\leq \frac{5^D}{4} \mathbb{E}_\varepsilon G \left(r \max_{\substack{v \in N_{\text{out}}, \\ m=1, \dots, k^{D-1}}} \max_{u \in N_{\text{in}} \cup \{b\}} \left| \sum_{i=1}^n \varepsilon_{i,v,m} x_{i,u} \right| \right). \end{aligned}$$

Proof of Theorem D.1. Without loss of generality, we assume that the parameters in Θ are rescaled with Algorithm 1, as the rescaling of parameters θ performed by Algorithm 1 does not change the associated function R_θ nor the path-norm $\|\Phi(\theta)\|_1$ so that we still have $\|\Phi(\theta)\|_1 \leq r$ and the supremum over $\theta \in \Theta$ on the left-hand side can be taken over rescaled parameters.

By induction on $d \geq 1$, we prove that

$$\begin{aligned} & \mathbb{E}_\varepsilon G \left(\sup_{\boldsymbol{\theta} \in \Theta} \sum_{\substack{i=1, \dots, n, \\ v \in N_{\text{out}}}} \varepsilon_{i,v} v(\boldsymbol{\theta}, x_i) \right) \\ & \leq \sum_{\ell=1}^d 5^{\ell-1} \mathbb{E}_\varepsilon G \left(r \max_{\substack{v \in N_{\text{out}}, \\ m=1, \dots, k^{\ell-1}}} \max_{u \in (\text{ant}^\ell(v) \cap N_{\text{in}}) \cup \{b\}} \left| \sum_{i=1}^n \varepsilon_{i,v,m} x_{i,u} \right| \right) \\ & \quad + \mathbb{E}_\varepsilon G \left(r \max_{v \in N_{\text{out}}} \max_{u \in \text{ant}^d(v) \setminus N_{\text{in}}} \sup_{\boldsymbol{\theta}} \left| \sum_{i=1}^n \varepsilon_{i,v} u(\boldsymbol{\theta}, x_i) \right| \right), \end{aligned}$$

with the convention that a maximum over an empty set is zero. This is true for $d = 1$ by Lemma D.2. The induction step is then verified using Lemma D.3. This concludes the induction. Applying the result for $d = D$, because $\text{ant}^D(v) \setminus N_{\text{in}} = \emptyset$, we get:

$$\begin{aligned} & \mathbb{E}_\varepsilon G \left(\sup_{\boldsymbol{\theta} \in \Theta} \sum_{\substack{i=1, \dots, n, \\ v \in N_{\text{out}}}} \varepsilon_{i,v} v(\boldsymbol{\theta}, x_i) \right) \\ & \leq \sum_{d=1}^D 5^{d-1} \mathbb{E}_\varepsilon G \left(r \max_{\substack{v \in N_{\text{out}}, \\ m=1, \dots, k^{d-1}}} \max_{u \in (\text{ant}^d(v) \cap N_{\text{in}}) \cup \{b\}} \left| \sum_{i=1}^n \varepsilon_{i,v,m} x_{i,u} \right| \right). \end{aligned}$$

We can only increase the right-hand side by considering maximum over all $u \in N_{\text{in}} \cup \{b\}$ and by adding independent copies indexed from $m = 1$ to $m = k^{D-1}$. Moreover, $\sum_{d=1}^D 5^{d-1} = (5^D - 1)/4$. This shows the final bound:

$$\begin{aligned} & \mathbb{E}_\varepsilon G \left(\sup_{\boldsymbol{\theta} \in \Theta} \sum_{\substack{i=1, \dots, n, \\ v \in N_{\text{out}}}} \varepsilon_{i,v} v(\boldsymbol{\theta}, x_i) \right) \\ & \leq \frac{5^D}{4} \mathbb{E}_\varepsilon G \left(r \max_{\substack{v \in N_{\text{out}}, \\ m=1, \dots, k^{D-1}}} \max_{u \in N_{\text{in}} \cup \{b\}} \left| \sum_{i=1}^n \varepsilon_{i,v,m} x_{i,u} \right| \right). \end{aligned}$$

□

E DETAILS TO DERIVE THE GENERALIZATION BOUND (THEOREM 3.1)

Proof of Theorem 3.1. Define the random matrices $E = (\varepsilon_{i,v})_{i,v} \in \mathbb{R}^{n \times d_{\text{out}}}$ and $R(\boldsymbol{\theta}, \mathbf{X}) = (v(\boldsymbol{\theta}, \mathbf{X}_i))_{i,v} \in \mathbb{R}^{n \times d_{\text{out}}}$ so that $\langle E, R(\boldsymbol{\theta}, \mathbf{X}) \rangle = \sum_{i,v} \varepsilon_{i,v} R_{\boldsymbol{\theta}}(\mathbf{X}_i)_v$. It then holds:

$$\begin{aligned} \mathbb{E}_{\mathbf{Z}} \ell\text{-generalization error of } \hat{\boldsymbol{\theta}}(\mathbf{Z}) & \leq \frac{2}{n} \mathbb{E}_{\mathbf{Z}, \varepsilon} \left(\sup_{\boldsymbol{\theta}} \sum_{i=1}^n \varepsilon_i \ell(R_{\boldsymbol{\theta}}(\mathbf{X}_i), \mathbf{Y}_i) \right) \\ & \leq \frac{2\sqrt{2}K}{n} \mathbb{E}_{\mathbf{Z}, \varepsilon} \left(\sup_{\boldsymbol{\theta}} \langle E, R(\boldsymbol{\theta}, \mathbf{X}) \rangle \right). \end{aligned}$$

The first inequality is the symmetrization property given by (Shalev-Shwartz & Ben-David, 2014, Theorem 26.3), and the second inequality is the vector-valued contraction property given by (Maurer, 2016). These are the relevant versions of very classical arguments that are widely used to reduce the problem to the Rademacher complexity of the model (Bach, Propositions 4.2 and 4.3)(Wainwright, 2019, Equations (4.17) and (4.18))(Bartlett & Mendelson, 2002, Proof of Theorem 8)(Shalev-Shwartz & Ben-David, 2014, Theorem 26.3)(Ledoux & Talagrand, 1991, Equation

(4.20)). In particular, this step has nothing specific with neural networks. Note that the assumption on the loss is used for the second inequality.

We now condition on $\mathbf{Z} = (\mathbf{X}, \mathbf{Y})$ and denote \mathbb{E}_ε the conditional expectation. For any random variable $\lambda(\mathbf{Z}) > 0$ measurable in \mathbf{Z} , it holds

$$\begin{aligned} \mathbb{E}_\varepsilon \left(\sup_{\boldsymbol{\theta}} \langle E, R(\boldsymbol{\theta}, \mathbf{X}) \rangle \right) &= \frac{1}{\lambda(\mathbf{Z})} \log \exp \left(\lambda(\mathbf{Z}) \mathbb{E}_\varepsilon \left(\sup_{\boldsymbol{\theta}} \langle E, R(\boldsymbol{\theta}, \mathbf{X}) \rangle \right) \right) \\ &= \frac{1}{\lambda(\mathbf{Z})} \log \exp \left(\mathbb{E}_\varepsilon \left(\lambda(\mathbf{Z}) \sup_{\boldsymbol{\theta}} \langle E, R(\boldsymbol{\theta}, \mathbf{X}) \rangle \right) \right) \\ &\leq \frac{1}{\lambda(\mathbf{Z})} \log \mathbb{E}_\varepsilon \exp \left(\lambda(\mathbf{Z}) \sup_{\boldsymbol{\theta}} \langle E, R(\boldsymbol{\theta}, \mathbf{X}) \rangle \right). \end{aligned}$$

For $z = ((x_i, y_i))_{i=1}^n \in (\mathbb{R}^{d_{\text{in}}} \times \mathbb{R}^{d_{\text{out}}})^n$, denote

$$e(z) = \mathbb{E}_\varepsilon \exp \left(\lambda(z) \sup_{\boldsymbol{\theta}} \langle E, R(\boldsymbol{\theta}, x) \rangle \right).$$

Since \mathbf{Z} is independent of ε , it holds

$$\mathbb{E}_\varepsilon \exp \left(\lambda(\mathbf{Z}) \sup_{\boldsymbol{\theta}} \langle E, R(\boldsymbol{\theta}, \mathbf{X}) \rangle \right) = e(\mathbf{Z}).$$

Denote $r = \sup_{\boldsymbol{\theta} \in \Theta} \|\Phi(\boldsymbol{\theta})\|_1$. For z as above, simply denote $\lambda := \lambda(z)$. The peeling argument given by Theorem D.1 for $G : t \in \mathbb{R} \mapsto \exp(\lambda t)$ gives:

$$e(z) \leq \frac{5^D}{4} \mathbb{E}_\varepsilon \exp \left(\lambda r \max_{m=1, \dots, k^{D-1}} \max_{v \in N_{\text{out}}, u \in N_{\text{in}} \cup \{b\}} \left| \sum_{i=1}^n \varepsilon_{i,v,m} x_{i,u} \right| \right),$$

where $x_{i,u}$ is coordinate u of vector $x_i \in \mathbb{R}^{d_{\text{in}}}$, and where b (for bias) is an added neuron for which we set by convention $x_b = 1$ for any input x . Denote

$$\sigma(x) := \max_{u \in N_{\text{in}} \cup \{b\}} \left(\sum_{i=1}^n x_{i,u}^2 \right)^{1/2} \geq \sqrt{n}.$$

Using Lemma E.1, it holds

$$\mathbb{E}_\varepsilon \exp \left(\lambda r \max_{\substack{v \in N_{\text{out}}, \\ u \in N_{\text{in}} \cup \{b\}, \\ m=1, \dots, k^{D-1}}} \left| \sum_{i=1}^n \varepsilon_{i,v,m}(\mathbf{X}_i)_u \right| \right) \leq 2k^{D-1}(d_{\text{in}} + 1)d_{\text{out}} \exp \left(\frac{(r\lambda(z)\sigma(x))^2}{2} \right).$$

Putting everything together, we get:

$$\mathbb{E}_\varepsilon \left(\sup_{\boldsymbol{\theta}} \langle E, R(\boldsymbol{\theta}, \mathbf{X}) \rangle \right) = e(\mathbf{Z}) \leq \left(\frac{1}{\lambda} \log(C_1) + \lambda(\mathbf{Z})C_2(\mathbf{X}) \right)$$

with

$$C_1 = 2k^{D-1}(d_{\text{in}} + 1)d_{\text{out}} \times \frac{5^D}{4} = \frac{5}{2}(5k)^{D-1}(d_{\text{in}} + 1)d_{\text{out}}$$

and

$$C_2(\mathbf{X}) = \frac{1}{2}(r\sigma(\mathbf{X}))^2.$$

Choosing $\lambda(\mathbf{Z}) = \sqrt{\frac{\log(C_1)}{C_2(\mathbf{X})}}$ yields:

$$\begin{aligned} \mathbb{E}_\varepsilon \left(\sup_{\boldsymbol{\theta}} \langle E, R(\boldsymbol{\theta}, \mathbf{X}) \rangle \right) &\leq 2\sqrt{\log(C_1)C_2(\mathbf{X})} \\ &\leq \underbrace{\sqrt{2}\sigma(\mathbf{X})r}_{=2\sqrt{C_2(\mathbf{X})}} \underbrace{\left(\log\left(\frac{5}{2}(d_{\text{in}} + 1)d_{\text{out}}\right) + D \log(5k) \right)^{1/2}}_{\geq \sqrt{\log(C_1)}}. \end{aligned}$$

Taking the expectation on both sides over \mathbf{Z} , and multiplying this by $\frac{2\sqrt{2}K}{n}$ yields Theorem 3.1. \square

The next lemma is classical (Golowich et al., 2018, Section 7.1) and is here only for completeness.

Lemma E.1. For any $d, k \in \mathbb{N}_{>0}$ and $\lambda > 0$, it holds

$$\mathbb{E}_{\epsilon} \exp \left(\lambda \max_{\substack{m=1, \dots, k, \\ u=1, \dots, d}} \left| \sum_{i=1}^n \epsilon_{i,m}(\mathbf{X}_i)_u \right| \right) \leq 2kd \max_{u=1, \dots, d} \exp \left(\frac{\lambda^2}{2} \sum_{i=1}^n (\mathbf{X}_i)_u^2 \right).$$

Proof. It holds

$$\mathbb{E}_{\epsilon} \exp \left(\lambda \max_{\substack{m=1, \dots, k, \\ u=1, \dots, d}} \left| \sum_{i=1}^n \epsilon_{i,m}(\mathbf{X}_i)_u \right| \right) \leq \sum_{\substack{m=1, \dots, k, \\ u=1, \dots, d}} \mathbb{E}_{\epsilon} \exp \left(\lambda \left| \sum_{i=1}^n \epsilon_{i,m}(\mathbf{X}_i)_u \right| \right).$$

For given u and m :

$$\begin{aligned} \mathbb{E}_{\epsilon} \exp \left(\lambda \left| \sum_{i=1}^n \epsilon_{i,m}(\mathbf{X}_i)_u \right| \right) &\stackrel{\text{Lemma D.1}}{\leq} 2 \mathbb{E}_{\epsilon} \exp \left(\lambda \sum_{i=1}^n \epsilon_{i,m}(\mathbf{X}_i)_u \right) \\ &= 2 \prod_{i=1}^n \frac{\exp(\lambda(\mathbf{X}_i)_u) + \exp(-\lambda(\mathbf{X}_i)_u)}{2} \leq 2 \exp \left(\frac{\lambda^2}{2} \sum_{i=1}^n (\mathbf{X}_i)_u^2 \right) \end{aligned}$$

using $\exp(x) + \exp(-x) \leq 2 \exp(x^2/2)$ in the last inequality. \square

F THE CROSS-ENTROPY LOSS IS LIPSCHITZ

Theorem 3.1 applies to the cross-entropy loss with $K = \sqrt{2}$. To see this, first recall that with C classes, the cross-entropy loss is defined as

$$\ell : (x, y) \in \mathbb{R}^C \times \{0, 1\}^C \rightarrow - \sum_{c=1}^{d_{\text{out}}} y_c \log \left(\frac{\exp(x_c)}{\sum_d \exp(x_d)} \right).$$

Consider $y \in \{0, 1\}^C$ with exactly one nonzero coordinate and an exponent $p \in [1, \infty]$ with conjugate exponent p' ($1/p + 1/p' = 1$). Then for every $x, x' \in \mathbb{R}^C$:

$$\ell(x, y) - \ell(x', y) \leq 2^{1/p'} \|x - x'\|_p.$$

Consider a class $c \in \{1, \dots, C\}$ and take $y \in \{0, 1\}^C$ to be a one-hot encoding of c (meaning that $y_{c'} = \mathbb{1}_{c'=c}$). Consider an exponent $p \in [1, \infty]$ with conjugate exponent p' ($1/p + 1/p' = 1$). The function $f : x \mapsto \ell(x, y) = - \sum_c y_c \log \left(\frac{\exp(x_c)}{\sum_{c'=1}^C \exp(x_{c'})} \right) = - \log \left(\frac{\exp(x_c)}{\sum_{c'=1}^C \exp(x_{c'})} \right)$ is continuously differentiable so that for every $x, x' \in \mathbb{R}^C$:

$$f(x) - f(x') = \int_0^1 \langle \nabla f(tx + (1-t)x'), x - x' \rangle dt \leq \sup_{t \in [0,1]} \|\nabla f(tx + (1-t)x')\|_p \|x - x'\|_{p'}.$$

In order to differentiate f , let's start to differentiate $g(x) = \frac{\exp(x_c)}{\sum_{c'=1}^C \exp(x_{c'})}$. Denote ∂_i the partial derivative with respect to coordinate i . Then for $i \neq c$:

$$\begin{aligned} \partial_c g(x) &= \frac{\exp(x_c) (\sum_{c'} \exp(x_{c'})) - \exp(x_c) (\exp(x_c))}{(\sum_{c'} \exp(x_{c'}))^2} \\ &= g(x) \frac{\sum_{c' \neq c} \exp(x_{c'})}{\sum_{c'} \exp(x_{c'})}. \\ \partial_i g(x) &= \frac{0 (\sum_{c'} \exp(x_{c'})) - \exp(x_c) (\exp(x_i))}{(\sum_{c'} \exp(x_{c'}))^2} \\ &= g(x) \frac{-\exp(x_i)}{\sum_{c'} \exp(x_{c'})}. \end{aligned}$$

Since $f(x) = (-\log \circ h)(x)$:

$$\begin{aligned}\partial_i f(x) &= -\frac{\partial_i g(x)}{g(x)} \\ &= \frac{1}{\sum_{c'=1}^C \exp(x_{c'})} \times \begin{cases} -\sum_{c' \neq c} e^{x_{c'}} & \text{if } i = c, \\ e^{x_i} & \text{otherwise.} \end{cases}\end{aligned}$$

Thus

$$\begin{aligned}\|\nabla f(x)\|_p^p &= \sum_{i=1}^C |\partial_i f(x)|^p \\ &= \frac{\left(\sum_{c' \neq c} \exp(x_{c'})\right)^p + \sum_{c' \neq c} \exp(x_{c'})^p}{\left(\sum_{c'=1}^C \exp(x_{c'})\right)^p} \\ &\leq 2 \frac{\left(\sum_{c' \neq c} \exp(x_{c'})\right)^p}{\left(\sum_{c'=1}^C \exp(x_{c'})\right)^p} \\ &\leq 2 \frac{\left(\sum_{c' \neq c} \exp(x_{c'})\right)^p}{\left(\sum_{c' \neq c} \exp(x_{c'})\right)^p} \\ &= 2.\end{aligned}$$

where we used in the first inequality that $\|v\|_p^p \leq \|v\|_1^p$ for any vector v . This shows that for every $x, x' \in \mathbb{R}^C$:

$$\ell(x, y) - \ell(x', y) \leq 2^{1/p} \|x - x'\|_{p'}.$$

G THE TOP-1 ACCURACY LOSS IS NOT LIPSCHITZ

Theorem 3.1 does not apply to the top-1 accuracy loss $\ell(\hat{y}, y) = \mathbb{1}_{\arg \max_{\hat{y}=\arg \max_y}$ as Equation (2) cannot be satisfied by ℓ . Indeed, it is easy to construct situations where $\hat{y}_1 = R_\theta(x_1)$ is arbitrarily close to $\hat{y}_2 = R_\theta(x_2)$ with x_2 correctly classified, while x_1 is not (just take x_2 on the boundary decision of the network and x_1 on the wrong side of the boundary), so that the left-hand side is equal to 1 and the right-hand side is arbitrarily small. There could thus not exist a finite $K > 0$ that satisfies Equation (2).

H THE MARGIN-LOSS IS LIPSCHITZ

For $\hat{y} \in \mathbb{R}^{d_{\text{out}}}$ and a one-hot encoding $y \in \mathbb{R}^{d_{\text{out}}}$ of the class c of x (meaning that $y_{c'} = \mathbb{1}_{c'=c}$ for every c'), the margin $M(\hat{y}, y)$ is defined by

$$M(\hat{y}, y) := [\hat{y}]_c - \max_{c' \neq c} [\hat{y}]_{c'}.$$

For $\gamma > 0$, recall that the γ -margin-loss is defined by

$$\ell(\hat{y}, y) = \begin{cases} 0 & \text{if } \gamma < M(\hat{y}, y), \\ 1 - \frac{M(\hat{y}, y)}{\gamma} & \text{if } 0 \leq M(\hat{y}, y) \leq \gamma, \\ 1 & \text{if } M(\hat{y}, y) < 0. \end{cases} \quad (13)$$

For any class c and one-hot encoding y of c , it is known that $\hat{y} \in \mathbb{R}^{d_{\text{out}}} \rightarrow M(\hat{y}, y)$ is 2-Lipschitz with respect to the L^2 -norm on \hat{y} (Bartlett et al., 2017, Lemma A.3). Moreover, the function

$$r \in \mathbb{R} \mapsto \begin{cases} 0 & \text{if } r < -\gamma, \\ 1 + \frac{r}{\gamma} & \text{if } -\gamma \leq r \leq 0, \\ 1 & \text{if } r > 0. \end{cases}$$

is $\frac{1}{\gamma}$ -Lipschitz. By composition, this shows that $\hat{y} \in \mathbb{R}^{d_{\text{out}}} \rightarrow \ell_\gamma(\hat{y}, y)$ is $\frac{2}{\gamma}$ -Lipschitz with respect to the L^2 -norm.

Proof of Theorem 3.2. Since the labels \mathbf{Y} are one-hot encodings, we equivalently consider \mathbf{Y} either in $\mathbb{R}^{d_{\text{out}}}$ or in $\{1, \dots, d_{\text{out}}\}$. It holds (Bartlett et al., 2017, Lemma A.4)

$$\mathbb{P}\left(\arg \max_c [R_{\theta}(\mathbf{X})]_c \neq \mathbf{Y}\right) \leq \mathbb{E}(\ell_{\gamma}(R_{\theta}(\mathbf{X}), \mathbf{Y}))$$

for any $\gamma > 0$ and associated γ -margin-loss ℓ_{γ} . Thus, considering the generalization error for ℓ_{γ} :

$$\mathbb{P}\left(\arg \max_c [R_{\theta}(\mathbf{X})]_c \neq \mathbf{Y}\right) \leq \underbrace{\frac{1}{n} \sum_{i=1}^n \ell_{\gamma}\left(R_{\hat{\theta}(\mathbf{Z})}(\mathbf{X}_i), \mathbf{Y}_i\right)}_{=\text{training error of } \hat{\theta}(\mathbf{Z})} + \mathbb{E}_{\mathbf{Z}} \ell_{\gamma}\text{-generalization error of } \hat{\theta}(\mathbf{Z}).$$

By definition of ℓ_{γ} , the training error of $\hat{\theta}(\mathbf{Z})$ is at most $\frac{1}{n} \sum_{i=1}^n \mathbb{1}_{[R_{\hat{\theta}(\mathbf{Z})}(\mathbf{X}_i)]_{\mathbf{Y}_i} \leq \gamma + \max_{c \neq \mathbf{Y}_i} [R_{\hat{\theta}(\mathbf{Z})}(\mathbf{X}_i)]_c}$. Moreover, Theorem 3.1 can be used to bound the generalization error associated with ℓ_{γ} with $K = 2/\gamma$. This proves the claim. \square

I DETAILS ON THE EXPERIMENTS OF SECTION 4

Details for Table 2. All the experiments are done on ImageNet-1k using 99% of the 1,281,167 images of the training set for training, the other 1% is used for validation. Thus, $n = 1268355 = \lfloor 0.99 \times 1281167 \rfloor$ in our experiments, $d_{\text{in}} = 224 \times 224 \times 3 = 150528$, $d_{\text{out}} = 1000$. We also estimated $B = 2.640000104904175$ by taking the maximum of the L^{∞} norms of the training images normalized for inference⁶. The PyTorch code for normalization at inference is standard:

```

1 inference_normalization = transforms.Compose([
2     transforms.Resize(256),
3     transforms.CenterCrop(224),
4     transforms.ToTensor(),
5     transforms.Normalize(mean=[0.485, 0.456, 0.406], std=[0.229,
6     0.224, 0.225]),
7 ])

```

We consider ResNets. They have a single max pooling layer of kernel size 3×3 so that $k = 9$. The depth is $D = 3 + \# \text{ basic blocks} \times \# \text{ conv per basic block}$, where 3 accounts for the conv1 layer, the average pooling layer, the fc layer, and the rest accounts for all the convolutional layers in the basic blocks. Table 4 details the relevant values related to basic blocks.

Table 4: Number of basic blocks, of convolutional layer per basic blocks and associated D_2 for ResNets (He et al., 2016, Table 1).

ResNet	18	34	50	101	152
# basic blocks	8	16		33	50
# conv per basic block	2		3		
D_2	18	34	50	101	152

Pretrained ResNets. The PyTorch pretrained weights that have been selected are the ones with the best performance: `ResNetX_Weights.IMAGENET1K_V1` for ResNets 18 and 34, and `ResNetX_Weights.IMAGENET1K_V2` otherwise.

Choice of $\gamma > 0$ for Theorem 3.2. In Equation (3), note that there is a trade-off when choosing $\gamma > 0$. Indeed, the first term of the right-hand side is non-decreasing with γ while the second

⁶The constant σ in Theorem 3.1 corresponds to data \mathbf{Z}_i drawn from the distribution for which we want to evaluate the test error. This is then the data normalized for inference. Thus, the training loss appearing in Theorem 3.1 is also evaluated on the training data \mathbf{Z}_i normalized for inference. In the experiments, we ignore this fact and still evaluate the training loss on the data augmented for training. Moreover, note that it is not possible to recover the training images augmented for training from the images normalized for inference, because cropping is done at random for training. Thus, the real life estimator is not a function of the images \mathbf{Z}_i normalized for inference, and Theorem 3.1 does not apply stricto sensu. This fact is ignored here.

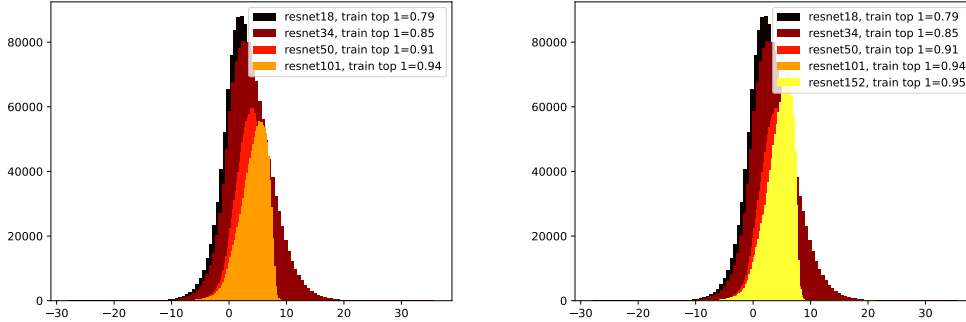


Figure 3: Distribution of the margins on the training set of ImageNet, with the pretrained ResNets available on PyTorch.

one is non-increasing. The first term is simply the proportion of datapoints that are not correctly classified with a margin at least equal to γ . Defining the margin of input i on parameters θ to be $R_{\theta}(\mathbf{X}_i)_{\mathbf{Y}_i} - \arg \max_{c \neq \mathbf{Y}_i} R_{\theta}(\mathbf{X}_i)_c$, this means that the first term is (approximately) equal to q if $\gamma = \gamma(q)$ is the q -quantile of the distribution of the margins over the training set.

Note that since the second term in Equation (3) is of order $1/\sqrt{n}$, it would be desirable to choose the $1/\sqrt{n}$ -quantile (up to a constant) for γ . However, this is not possible in practice as soon as the training top 1 accuracy is too large compared to $1/\sqrt{n}$ (eg. on ImageNet). Indeed, if the training top 1 error is equal to $e \in [0, 1]$, then at least a proportion e of the data margins should be negative⁷ so that any q -quantile with $q < e$ is negative and cannot be considered for Theorem 3.2

The distribution of the margins on the training set of ImageNet can be found in Figure 3. The maximum training margin is roughly of size 30, which is insufficient to compensate the size of the L^1 path-norm of pretrained ResNets reported in Table 3. For $\gamma > 30$, the first term of the right hand-side of Theorem 3.2 is greater than one, so that the bound is not informative. This shows that there is no possible choice for $\gamma > 0$ that makes the bound informative on these pretrained ResNets. Table 5 reports a quantile for these pretrained ResNets.

Table 5: The q -quantile $\gamma(q)$ for $q = \frac{1}{3}e + \frac{2}{3}$, with e being the top 1 error, on ImageNet, of pretrained ResNets available on PyTorch.

ResNet	18	34	50	101	152
$\gamma(q)$	5.0	5.6	4.2	5.6	5.8

Details for sparse networks. ResNet18 is trained on 99% of ImageNet with a single GPU using SGD for 90 epochs, learning rate 0.1, weight-decay 0.0001, batch size 1024, and a multi-step scheduler where the learning rate is divided by 10 at epochs 30, 60 and 80. The epoch out of the 90 ones with maximum validation top 1 accuracy is considered as the final epoch. Pruning is done iteratively according to (Frankle et al., 2021). We prune 20% of the remaining weights of each convolutional layer, and 10% of the final fully connected layer, at each pruning iteration, save the mask and rewind the weights to their values after the first 5 epochs of the dense network, and train for 85 remaining epochs, before pruning again etc. Results for a single run are shown in Figure 4.

Details for increasing the train size. Instead of training on 99% of ImageNet ($n = 1268355$), we trained a ResNet18 on $n/2^k$ samples drawn at random, for $1 \leq k \leq 5$. For each given k , the results are averaged over 3 seeds. The hyperparameters are the same as for sparse networks (except that we do not perform any pruning here): 90 epochs etc. Results are in Figure 5.

⁷A data margin is negative if and only if it is misclassified.

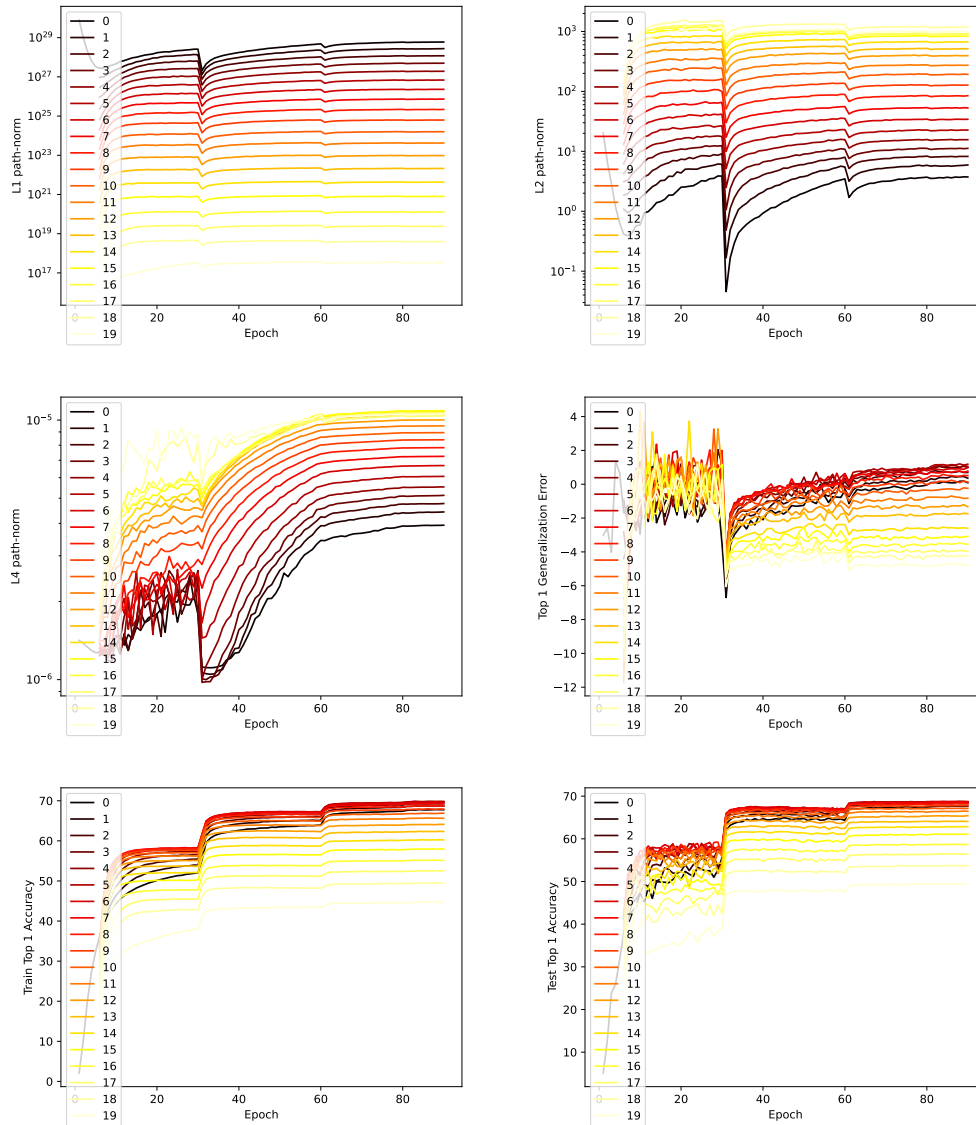


Figure 4: L^q path-norm ($q = 1, 2, 4$), test top 1 accuracy, training top 1 accuracy, and the top 1 generalization error (difference between test top 1 and train top 1) during the training of a ResNet18 on ImageNet. The pruning iteration is indicated in legend, with 0 corresponding to the dense network. The color also indicates the degree of sparsity: from dense (black) to extremely sparse (yellow).




Figure 5: L^1 path-norm, and empirical generalization errors for both the top-1 accuracy and the cross-entropy during the training of a ResNet18 on a subset of the training images of ImageNet. The legend indicates the size of the subset considered, *e.g.* $1/m$ corresponds to $1/m$ of 99% of ImageNet, leaving the other 1% out for validation. The color also indicates the size of the subset: from small (black) to large (yellow).

J MAIN RELATED WORKS

Given the extent of the literature on generalization bounds, we apologize in advance for papers the reader may find missing below.

Previous definitions of the path-embedding and the path-activations In the work of (Kawaguchi et al., 2017, Section 5.1) the path-embedding and the path-activations are evoked in the case of a ReLU DAG with max pooling neurons *and no biases*, but with no explicit definitions. Definition A.1 gives a formal definition for these objects, and extend it to the case where there are biases, which requires extending the path-embedding to paths starting from neurons that are not input neurons. Note that the formula Equation (1) is already stated in (Kawaguchi et al., 2017, Section 5.1) (as an explicit sum rather than an inner product), without proof since the objects are not explicitly defined in Kawaguchi et al. (2017).

A formal definition of the path-embedding is given in the specific case of ReLU feedforward neural networks with biases in the work of (Stock & Gribonval, 2022, Definition 6). Moreover, it is proved that Equation (1) holds *in this specific case* in (Stock & Gribonval, 2022, Corollary 3). Definition A.1 and Equation (1) generalize the latter to an arbitrary DAG with max or identity neurons (allowing for skip connections, max and average pooling).

The rest of the works we are aware of only define and consider the norm of the path-embedding, but not the embedding itself. The most general setting being the one of (Neyshabur et al., 2015) with a general DAG, *but without max or identity neurons, nor biases*. Not defining the path-embedding and the path-activations makes notations arguably heavier since Equation (1) is then always written with an explicit sum over all paths, with explicit product of weights along each path, and so on.

Previous generalization bounds based on path-norm See Table 1 for a comparison. Appendix K tackles some other bounds that do not appear in Table 1.

Empirical evaluation of path-norm Formula given in Theorem A.1 is the first one to fully encompass modern networks with biases, max/average pooling, and skip connections, such as ResNets. An equivalent formula is stated for ReLU *feedforward* networks *without biases* (and *no pooling/skip connections*) in (Dziugaite et al., 2020, Appendix C.6.5) and (Jiang et al., 2020, Equations (43) and (44)) but without proof. Actually, *this equivalent formula turns out to be false when there are max pooling neurons* as one must replace max pooling neurons with identity ones, see Theorem A.1. Care must also be taken with average pooling neurons that must be rescaled by considering them as identity neurons.

We could not find reported numerical values of the path-norm except for toy examples (Dziugaite, 2018; Furusho, 2020; Zheng et al., 2019). Details are in Appendix K.

Appendix K also discusses 1) inherent limitations of Theorem 3.1 which are common to many generalization bounds, and 2) the applicability of Theorem 3.1 compared to PAC-Bayes bounds.

K MORE RELATED WORKS

More generalization bounds using path-norm (E et al., 2022) establishes an additional bound to the one appearing in Table 1, for a class of functions and a complexity measure that are related to the *infinite-depth* limits of residual networks and the path-norm. However, it is unclear how this result implies anything for actual neural networks with the actual path-norm ⁸.

The bound in (Zheng et al., 2019) holds only for ReLU feedforward neural networks (*no max or identity neurons*) with *no biases* and *it grows exponentially* with the depth of the network. It is not included in Table 1 because it requires an additional assumption: the coordinates of the path-embedding must not only be bounded from above, but also *from below*. The reason for this assump-

⁸(E et al., 2022) starts from the fact that the infinite-depth limits of residual networks can be characterized with partial differential equations. Then, (E et al., 2022) establishes a bound for functions characterized by similar, but different, partial differential equations, using what seems to be an analogue of path-norm for these new functions. However, even if the characterizations of these functions are closed, as it is said in (E et al., 2022), "it is unclear how the two spaces are related".

tion is not discussed in (Zheng et al., 2019), and it is unclear whether this is at all desirable, since a small path-norm can only be better for generalization in light of Theorem 3.1.

(Golowich et al., 2018, Theorem 4.3) is a bound that holds for ReLU *feedforward* neural networks *with no biases (no max and no identity neurons)* and it *depends on the product of operators' norms of the layers*. It has the merit of having no dependence on the size of the architecture (depth, width, number of neurons etc.). However, it requires an additional assumption: each layer must have an operator norm bounded from below, so that it only applies to a restricted set of feedforward networks. Moreover, it is unclear whether such an assumption is desirable: there are networks with arbitrary small operators' norms that realize the zero function, and the latter has a generalization error equal to zero.

Theorem 8 in (Kawaguchi et al., 2017) gives a generalization bound for *scalar-valued* ($d_{\text{out}} = 1$) models with an output of the form $\langle \Phi(\theta), \mathbf{A}(\theta', x)x \rangle$ for some specific parameters θ, θ' that have no reason to be equal. This is orthogonal to the case of neural networks where one must have $\theta = \theta'$, and it is therefore not included in Table 1. Theorem 5 in (Kawaguchi et al., 2017) can be seen as a possible first step to derive a bound based on path-norm in the specific case of the mean squared error loss. However, as discussed in more details below, (Kawaguchi et al., 2017, Theorem 5) is a rewriting of the generalization error with several terms that are as complex to bound as the original generalization error, resulting in a bound being as hard as the generalization error to evaluate/estimate.

More details about Theorem 5 in (Kawaguchi et al., 2017) We start by re-deriving Theorem 5 in (Kawaguchi et al., 2017). In the specific case of mean squared error, using that

$$\|R_{\theta}(x) - y\|_2^2 = \|R_{\theta}(x)\|_2^2 + \|y\|_2^2 - 2 \langle R_{\theta}(x), y \rangle,$$

it is possible to rewrite the generalization error as follows:

$$\begin{aligned} \text{generalization error of } \hat{\theta}(\mathbf{Z}) &= \mathbb{E} \left(\|R_{\hat{\theta}(\mathbf{Z})}(\tilde{\mathbf{X}}) - \tilde{\mathbf{Y}}\|_2^2 | \mathbf{Z} \right) - \frac{1}{n} \sum_{i=1}^n \|R_{\hat{\theta}(\mathbf{Z})}(\mathbf{X}_i) - \mathbf{Y}_i\|_2^2 \\ &= \mathbb{E} \left(\|R_{\hat{\theta}(\mathbf{Z})}(\tilde{\mathbf{X}})\|_2^2 | \mathbf{Z} \right) - \frac{1}{n} \sum_{i=1}^n \|R_{\hat{\theta}(\mathbf{Z})}(\mathbf{X}_i)\|_2^2 \\ &\quad + \mathbb{E} \left(\|\tilde{\mathbf{Y}}\|_2^2 | \mathbf{Z} \right) - \sum_{i=1}^n \|\mathbf{Y}_i\|_2^2 \\ &\quad - 2\mathbb{E} \left(\left\langle R_{\hat{\theta}(\mathbf{Z})}(\tilde{\mathbf{X}}), \tilde{\mathbf{Y}} \right\rangle | \mathbf{Z} \right) - \frac{1}{n} \sum_{i=1}^n \left\langle R_{\hat{\theta}(\mathbf{Z})}(\mathbf{X}_i), \mathbf{Y}_i \right\rangle. \end{aligned}$$

It is then possible to make the L^2 path-norm appear. For instance, for one-dimensional output networks, it can be proven (see Lemma A.1) that $R_{\theta}(x) = \langle \Phi(\theta), z(x, \theta) \rangle$ with $\Phi(\theta)$ the path-embedding of parameters θ , and z that typically depends on the path-activations and the input, so that the first term above can be rewritten

$$\Phi(\hat{\theta}(\mathbf{Z}))^T \left(\mathbb{E} \left(z(\tilde{\mathbf{X}}, \hat{\theta}(\mathbf{Z})) z(\tilde{\mathbf{X}}, \hat{\theta}(\mathbf{Z}))^T | \mathbf{Z} \right) - \frac{1}{n} \sum_{i=1}^n z(\mathbf{X}_i, \hat{\theta}(\mathbf{Z})) z(\mathbf{X}_i, \hat{\theta}(\mathbf{Z}))^T \right) \Phi(\hat{\theta}(\mathbf{Z})).$$

Let us call a "generalization error like quantity" any term of the form

$$\mathbb{E} \left(f_{\hat{\theta}(\mathbf{Z})}(\tilde{\mathbf{Z}}) | \mathbf{Z} \right) - \frac{1}{n} \sum_{i=1}^n f_{\hat{\theta}(\mathbf{Z})}(\mathbf{Z}_i),$$

that is, any term that can be represented as a difference between the estimator learned from training data \mathbf{Z} evaluated on test data $\tilde{\mathbf{Z}}$, and the evaluation on the training data. We see that the derivation above replaces the classical generalization error to two others quantities similar in definition to the generalization error. This derivation, which is specific to mean squared error, leads to Theorem 5 in (Kawaguchi et al., 2017). Very importantly, note that this derivation trades a single quantity similar to generalization error for two new such quantities. (Kawaguchi et al., 2017) does not discuss how to bound these two terms, but without any further new idea, there is no other way than the ones

developed in the literature so far: reduce the problem to bounding a Rademacher complexity (as it is done in Theorem 3.1), or use the PAC-Bayes framework, and so on.

More on numerical evaluation of path-norm (Dziugaite, 2018, Section 2.9.1) reports numerical evaluations after 5 epochs of SGD on a one hidden layer network trained on a binary variant of MNIST. (Furusho, 2020, Figure 9 and Section 3.3.1) deals with 1d regression with 5 layers and 100 width. (Zheng et al., 2019) experiments on MNIST. Note that it is not clear whether (Zheng et al., 2019) reports the path-norm as defined in Definition A.1. Indeed, (Zheng et al., 2019) quotes both (Neyshabur et al., 2015) and (Neyshabur et al., 2017) when referring to the path-norm, but these two papers have two different definitions of the path-norm, as (Neyshabur et al., 2017) normalize it by the margin while (Neyshabur et al., 2015) does not.

For completeness, let us also mention that (Dziugaite et al., 2020; Jiang et al., 2020) reports whether the path-norm correlates with the empirical generalization error or not, but do not report the numerical values. (Neyshabur et al., 2017) reports the path-norm normalized by the margins, but not separately from each other.

Inherent limitations of uniform convergence bounds Theorem 3.1 has some inherent limitations due to its nature. It is *data-dependent* as it depends on the input distribution. However, it does not depend on the label distribution, making it uninformative as soon as Θ is so much expressive that it can fit random labels. Networks that can fit random labels have already been found empirically (Zhang et al., 2021), and it is open whether this stays true with a constraint on the path-norm.

Theorem 3.1 is based on a uniform convergence bound⁹ as any other bound also based on a control of a Rademacher complexity. (Nagarajan & Kolter, 2019) empirically argue that even the tightest uniform convergence bound holding with high probability must be loose on some synthetic datasets. If this was confirmed theoretically, this would still allow uniform bounds to be tight when considering other datasets than the one in (Nagarajan & Kolter, 2019), such as real-world datasets, or when the estimator considered in (Nagarajan & Kolter, 2019) is not in Θ (for instance because of constraints on the slopes via the path-norm).

Finally, Theorem 3.1 can provide theoretical guarantees on the generalization error of the output of a learning algorithm, but only *a posteriori*, after training. In order to have *a priori* guarantees, one should have to derive *a priori* knowledge on the path-norm at the end of the learning algorithm.

Comparison to PAC-Bayes bounds Another interesting direction to understand generalization of neural networks is the PAC-Bayes framework (Guedj, 2019; Alquier, 2021). Unfortunately, PAC-Bayes bounds cannot be exactly computed on networks that are trained in a usual way. Indeed, these bounds typically involve a KL-divergence, or something related, for which there is no closed form except for very specific distributions (iid Gaussian/Cauchy weights...) that do not correspond to the distributions of the weights after a usual training^{10,11}. We are aware of two research directions that try to get over this issue. The first way is to change the learning algorithm by enforcing the weights to be iid normally distributed, and then optimize the parameters of these normal distributions, see for instance the pioneer work (Dziugaite & Roy, 2017). The merit of this new learning algorithm is that it has explicitly been designed with the goal of having a small generalization error. Practical performance are worse than with usual training, but this leads to networks with an associated non-vacuous generalization bound¹², and unfortunately, this does not apply to usual training. The second way to get over the intractable evaluation of the KL-divergence is to 1) try to approximate the bound within reasonable time, and 2) try to quantify the error made with the approximation (Pérez & Louis, 2020). Unfortunately, to the best of our knowledge, approximation is often based on a distribution assumption of the weights that is not met in practice (*e.g.* iid Gaussian weights), approximation is costly, and the error is unclear when applied to networks trained usually. For instance, the bound in (Pérez & Louis, 2020, Section 5) 1) requires at least $O(n^2)$ operations to be evaluated, with n being the number of training examples, thus being prohibitive for large n (Pérez & Louis, 2020, Section

⁹A uniform convergence bound on a model class \mathcal{F} is a bound on $\mathbb{E}_{\mathcal{Z}} \sup_{f \in \mathcal{F}} f(Z)$ generalization error $f(Z)$. This worst-case type of bound can lead to potential limitations when \mathcal{F} is too expressive.

¹⁰The randomness of the weights after training comes from the random initialization and the randomness in the algorithm (*e.g.* random batch in SGD).

¹¹For instance, independence is not empirically observed, see (Frankle et al., 2020, Section 5.2)

¹²Except, of course, for methods that are based on the evaluation of the performance on held-out data.

7), and 2) it is unclear what error is being made when using a Gaussian process as an approximation of the neural network learned by SGD.