
Approximating Lipschitz continuous functions with GroupSort neural networks

U. Tanielian
Criteo, Sorbonne Université
u.tanielian@criteo.com
Paris, France

M. Sangnier
Sorbonne Université
maxime.sangnier@upmc.fr
Paris, France

G. Biau
Sorbonne Université
gerard.biau@upmc.fr
Paris, France

Abstract

Recent advances in adversarial attacks and Wasserstein GANs have advocated for use of neural networks with restricted Lipschitz constants. Motivated by these observations, we study the recently introduced GroupSort neural networks, with constraints on the weights, and make a theoretical step towards a better understanding of their expressive power. We show in particular how these networks can represent any Lipschitz continuous piecewise linear functions. We also prove that they are well-suited for approximating Lipschitz continuous functions and exhibit upper bounds on both the depth and size. To conclude, the efficiency of GroupSort networks compared with more standard ReLU networks is illustrated in a set of synthetic experiments.

1 Introduction

In the past few years, developments in deep learning have highlighted the benefits of operating neural networks with restricted Lipschitz constants. An important illustration is provided by robust machine learning, where networks with large Lipschitz constants are prone to be more sensitive to adversarial attacks, in the sense that small perturbations of the inputs can lead to significant misclassification errors (e.g., [Goodfellow et al., 2015](#)). In order to circumvent these limitations, [Gao et al. \(2017\)](#), [Esfahani and Kuhn \(2018\)](#), and [Blanchet et al. \(2019\)](#) studied a new regularization scheme based on penalizing the gradients of the networks. Constrained neural networks also play a key role in the different but not less important domain of Wasserstein GANs ([Arjovsky et al., 2017](#)), which take advantage

of the dual form of the 1-Wasserstein distance expressed as a supremum over the set of 1-Lipschitz functions ([Villani, 2008](#)). This formulation has been shown to bring training stability and is empirically efficient ([Gulrajani et al., 2017](#)). In this context, many different ways have been explored to restrict the Lipschitz constants of the discriminator. One possibility is to clip their weights, as advocated by [Arjovsky et al. \(2017\)](#). Other solutions involve enforcing a gradient penalty ([Gulrajani et al., 2017](#)) or penalizing norms of the matrices of the weights ([Miyato et al., 2018](#)).

However, all of these operations are delicate and may significantly affect the expressive power of the neural networks. For example, [Huster et al. \(2018\)](#) show that ReLU neural networks with constraints on the weights cannot represent even the simplest functions, such as the absolute value. In fact, little is known regarding the expressive power of such restricted networks, since most studies interested in the expressiveness of neural networks (e.g., [Hornik et al., 1989](#); [Cybenko, 1989](#); [Raghu et al., 2017](#)) do not take into account eventual constraints on their architectures. As far as we know, the most recent attempt to tackle this issue is by [Anil et al. \(2019\)](#). These authors exhibit a family of neural networks, with constraints on the weights, which is dense in the set of Lipschitz continuous functions on a compact set. To show this result, [Anil et al. \(2019\)](#) make critical use of GroupSort activations.

Motivated by the above, our objective in the present article is to make a step towards a better mathematical understanding of the approximation properties of Lipschitz feedforward neural networks using GroupSort activations. Our contributions are threefold:

- (i) We show that GroupSort neural networks, with constraints on the weights, can represent any Lipschitz continuous piecewise linear function and exhibit upper bounds on both their depth and size. We make a connection with the literature on the depth and size of ReLU networks (in particular [Arora et al., 2018](#); [He et al., 2018](#)).
- (ii) Building on the work of [Anil et al. \(2019\)](#), we offer

upper bounds on the depth and size of GroupSort neural networks that approximate 1-Lipschitz continuous functions on compact sets. We also show that increasing the grouping size may significantly improve the expressivity of GroupSort networks.

- (iii) We empirically compare the performances of GroupSort and ReLU networks in the context of function regression estimation and Wasserstein distance approximation.

The mathematical framework together with the necessary notation is provided in Section 2. Section 3 is devoted to the problem of representing Lipschitz continuous functions with GroupSort networks of grouping size 2. The extension to any arbitrary grouping size is discussed in Section 4 and numerical illustrations are given in Section 5. For the sake of clarity, all proofs are gathered in the Appendix.

2 Mathematical context

We introduce in this section the mathematical context of the article and describe more specifically the GroupSort neural networks, which, as we will see, play a key role in representing and approximating Lipschitz continuous functions.

Throughout the paper, the ambient space \mathbb{R}^d is assumed to be equipped with the Euclidean norm $\|\cdot\|_k$. For E a subset of \mathbb{R}^d , we denote by $\text{Lip}_1(E)$ the set of 1-Lipschitz real-valued functions on E , i.e.,

$$\text{Lip}_1(E) = \{f : E \rightarrow \mathbb{R} : |f(x) - f(y)| \leq \|x - y\|_k, (x, y) \in E^2\}$$

Let $k > 2$ be an integer. We let $D_k = \{D_{k;a} : a \in L\}$ be the class of functions from \mathbb{R}^d to \mathbb{R} parameterized by feedforward neural networks of the form

$$D_{k;a}(x) = \begin{matrix} V_q & S_k(& V_q & 1 & & S_k(& V_2 & S_k(& V_1 & x + & c_1 &) \\ 1 & v_{q-1} & v_{q-1} & v_{q-2} & & v_2 & v_1 & v_1 & D & v_1 & 1 & \\ & & & & & & & & & & & \end{matrix} + \begin{matrix} c_2 \\ v_2 & 1 \end{matrix} + \begin{matrix} c_q & 1 \\ v_q & 1 \end{matrix} + \begin{matrix} c_q \\ 1 & 1 \end{matrix} \quad (1)$$

where $q > 2$ and the characters below the matrices indicate their dimensions (lines – columns). For $q = 1$, we simply let $D_{k;a}(x) = V_1 x + c_1$ be a simple linear regression in \mathbb{R} without hidden layers. Thus, a network in D_k has $(q - 1)$ hidden layers, and hidden layers from depth 1 to $(q - 1)$ are assumed to be of respective widths $v_i, i = 1, \dots, q - 1$, divisible by k . Such a network is said to be of depth q and of size $n_1 + \dots + n_{q-1}$. The matrices V_i are the matrices of weights between layer i and layer $(i + 1)$ and the c_i 's are the corresponding offset vectors (in column format). So, altogether, the vectors $a = (V_1, \dots, V_q; c_1, \dots, c_q)$ represent the parameter space L of the functions in D_k . With respect to the activation functions S_k , we propose to use the GroupSort activation, which separates the pre-activations into groups and then sorts each group into ascending order.




Figure 1: GroupSort activation with a grouping size 5. Source: Anil et al. (2019).

The GroupSort function splits the input into n different groups of k elements each: $G_1 = \{x_1, \dots, x_k\}, \dots, G_n = \{x_{nk-(k-1)}, \dots, x_{nk}\}$, and then orders each group by decreasing order. Thus, the GroupSort function with a grouping size $k > 2$ is applied on a given vector (x_1, \dots, x_{kn}) as follows:

$$S_k(x_1, \dots, x_k, \dots, x_{nk-(k-1)}, \dots, x_{nk}) = (x_{(k)}^{G_1}, \dots, x_{(1)}^{G_1}, \dots, x_{(k)}^{G_n}, \dots, x_{(1)}^{G_n});$$

where $(x_{(k)}^{G_i}, \dots, x_{(1)}^{G_i})$ corresponds to the decreasing ordering in the group G_i .

This activation is applied on groups of k components, which makes sense in (1) since the widths of the hidden layers are assumed to be divisible by k . GroupSort has been introduced in Anil et al. (2019) as a 1-Lipschitz activation function that preserves the gradient norm of the input. An example with a grouping size $k = 5$ is given in Figure 1. With a slight abuse of vocabulary, we call a neural network of the form (1) a GroupSort neural network. We note that the GroupSort activation can recover the standard rectifier function. For example, $S_2(x; 0) = (\text{ReLU}(x); \text{ReLU}(-x))$, but the converse is not true.

Throughout the manuscript, the notation $\|\cdot\|_k$ (respectively, $\|\cdot\|_{k_\infty}$) means the Euclidean (respectively, the supremum) norm on \mathbb{R}^p , with no reference to p as the context is clear. For $W = (w_{i,j})$ a matrix of size $p_1 \times p_2$, we let $\|W\|_{k_2} = \sup_{\|x\|_k=1} \|Wx\|_k$ be the 2-norm of W . Similarly, the ∞ -norm of W is $\|W\|_{k_\infty} = \sup_{\|x\|_{k_\infty}=1} \|Wx\|_{k_\infty} = \max_{i=1, \dots, p_1} \sum_{j=1}^{p_2} |w_{i,j}|$. We will also use the $(2; \infty)$ -norm of W , i.e., $\|W\|_{k_2, \infty} = \sup_{\|x\|_k=1} \|Wx\|_{k_\infty}$. The following assumption plays a central role in our approach:

Assumption 1. For all $a = (V_1, \dots, V_q; c_1, \dots, c_q) \in L$,

$$\|V_1\|_{k_2, \infty} \leq 1; \max(\|V_2\|_{k_\infty}, \dots, \|V_q\|_{k_\infty}) \leq 1; \\ \text{and } \max(\|c_i\|_{k_\infty} : i = 1, \dots, q) \leq K_2;$$

where $K_2 > 0$ is a constant.

This type of compactness requirement has already been suggested in the statistical and machine learning community (e.g., Arjovsky et al., 2017; Anil et al., 2019; Biau et al., 2020). In the setting of this article, its usefulness is captured in the following simple but essential lemma:

Lemma 1. Assume that Assumption 1 is satisfied. Then, for any $k > 2$, $D_k \subset \text{Lip}_1(\mathbb{R}^d)$.

Combining Lemma 1 with Arzelà-Ascoli theorem, it is easy to see that, under Assumption 1, the class \mathcal{D}_K restricted to any compact $K \subset \mathbb{R}^d$ is compact in the set of continuous functions on K with respect to the uniform norm. From this point of view, Assumption 1 is therefore somewhat restrictive. On the other hand, it is essential in order to guarantee that all neural networks in \mathcal{D}_K are indeed 1-Lipschitz. Practically speaking, various approaches have been explored in the literature to enforce this Lipschitz constraint. Gulrajani et al. (2017), Kodali et al. (2017), Wei et al. (2018), and Zhou et al. (2019) proposed a gradient penalty term, Miyato et al. (2018) applied spectral normalization, while Anil et al. (2019) have shown the empirical efficiency of the orthonormalization of Björck and Bowie (1971).

Importantly, Anil et al. (2019, Theorem 3) states that, under Assumption 1, GroupSort neural networks are universal Lipschitz approximators on compact sets. More precisely, for any Lipschitz continuous function defined on a compact, one can find a neural network of the form (1) verifying Assumption 1 and arbitrarily close to f with respect to the uniform norm. Our objective in the present article is to explore the properties of these networks. We start in the next section by examining the case of piecewise linear functions.

3 Learning functions with a grouping size 2

For this section, we only consider GroupSort neural networks with a grouping size 2 and aim at studying their expressivity. The capacity of GroupSort networks to approximate continuous functions is studied via the representation of piecewise linear functions. For feedforward ReLU networks, their ability to represent such functions has been largely studied. In particular, Arora et al. (2018, Theorem 2.1) reveals that any piecewise linear function from \mathbb{R}^d to \mathbb{R} can be represented by a ReLU network of depth at most $\log_2(d+1)e$ (the symbol e stands for the ceiling function), whereas He et al. (2018) specify an upper bound on their size. In the present section, we extend these results and first tackle the problem of representing piecewise linear functions with constrained GroupSort networks. Then we move to the non-linear case.

3.1 Representation of piecewise linear functions

Let us start gently by fixing the vocabulary.

Definition 1. A continuous function $f: \mathbb{R}^d \rightarrow \mathbb{R}$ is said to be (continuous) m_f -piecewise linear (if $m_f > 2$) if there exist a partition $W = \{W_1, \dots, W_{m_f}\}$ of \mathbb{R}^d into polytopes and a collection $\{\phi_1, \dots, \phi_{m_f}\}$ of affine functions such that, for all $x \in W_i$, $i = 1, \dots, m_f$, $f(x) = \phi_i(x)$.

At this stage no further assumption is made on the sets W_1, \dots, W_{m_f} , which are just assumed to be polytopes. An example of piecewise linear function on the real line with $m_f = 4$ is depicted in Figure 2. As this figure suggests, $\min(f_1, \dots, f_m)$.

Figure 2: A 4-piecewise linear function on the real line and the associated partition $W = \{W_1, \dots, W_4\}$ and $\tilde{W} = \{\tilde{W}_1, \dots, \tilde{W}_7\}$. The partition \tilde{W} is finer than W .

the ambient space \mathbb{R}^d can be further covered by a second partition $\tilde{W} = \{\tilde{W}_1, \dots, \tilde{W}_{M_f}\}$ of M_f polytopes ($M_f > 1$), in such a way that the sign of the differences $\phi_j - \phi_i$, $(i, j) \in \{1, \dots, m_f\}^2$, does not change on the subsets $\tilde{W}_1, \dots, \tilde{W}_{M_f}$. It is easy to see that the partition \tilde{W} is finer than W since, for each $i \in \{1, \dots, m_f\}$ there exists $j \in \{1, \dots, m_f\}$ such that $\tilde{W}_j \subset W_i$. This implies in particular that $M_f > m_f$.

The usefulness of the partition \tilde{W} is demonstrated by He et al. (2018, Theorem 5.1), which states that any piecewise linear function f can be written as

$$f = \max_{k \in \{1, \dots, m_f\}} \min_{i \in S_k} \phi_i, \quad (2)$$

where each S_k is a non-empty subset of $\{1, \dots, m_f\}$. This characterization of the function f is interesting, since it shows that any m_f -piecewise linear function can be computed using only a finite number of max and min operations. As identity (2) is essential for our approach, this justifies spending some time examining it.

Lemma 2. Let $f: \mathbb{R}^d \rightarrow \mathbb{R}$ be an m_f -piecewise linear function. Then $m_f \leq M_f \leq \min(2^{m_f-2}, (m_f - 2)^{2d})$.

Lemma 2 is an improvement of He et al. (2018, Lemma 5.1), which shows that $M_f \leq m_f!$. Our proof method exploits the inequality $M_f \leq C_{m_f(m_f-1)=2,d}$, where $C_{n,d}$ denotes the number of arrangements of hyperplanes in a space of dimension d (Devroye et al., 1996, Chapter 5). Another application of (2) is encapsulated in Proposition 1 below, which will be useful for later analysis, in combining maxima and minima in neural networks of the form (1).

Proposition 1. Let $f_1, \dots, f_m: \mathbb{R}^d \rightarrow \mathbb{R}$ be a collection of functions ($m > 2$), each represented by a neural network of the form (1), with common depth q and sizes s_i , $i = 1, \dots, m$.

In the specific case where $m = 2^n$ for some $n > 1$, there exist neural networks of the form (1) (with grouping size 2) with depth $q + \log_2(m)$ and size at most $s_1 + \dots + s_m + m - 1$ that represent the function $f = \max(f_1, \dots, f_m)$ and $g = \min(f_1, \dots, f_m)$.

If m is arbitrary, then there exist neural networks of the form (1) with depth $q + d \log_2(m)$ and size at most $2m^{q+1} + s_m + 1$ that represent the functions f and g .

Interestingly, Arora et al. (2018, Lemma D.3), which is the analog of Proposition 1 asserts that the size with ReLU activations is at most $s_1 + s_m + 8m^4$. For the specific computation of maxima/minima of functions, it should be stressed that GroupSort activations slightly reduces the size of the networks. By combining Lemma 2, Proposition 1 and identity (2), we are led to the following theorem, which reveals the ability of GroupSort networks for representing 1-Lipschitz piecewise linear functions.

Theorem 1. Let $f \in \text{Lip}_1(\mathbb{R}^d)$ that is also m_f -piecewise linear. Then there exists a neural network of the form (1) verifying Assumption 1 that represents f . Besides, its depth is $d \log_2(M_f) + d \log_2(m_f) + 1$ and its size is at most $3m_f M_f + M_f - 1$.

This result should be compared with state-of-the-art results known for ReLU neural networks. In particular, Arora et al. (2018, Theorem 2.1) reveals that any m_f -piecewise linear function f can be represented by a ReLU network with depth at most $d \log_2(d + 1)e$. The upper bound of Theorem 1 can be larger since it involves both M_f and m_f . On the other hand, the upper bound $O(m_f M_f)$ on the size significantly improves on He et al. (2018, Theorem 5.2), which is at least $O(d 2^{m_f M_f})$. This improvement in terms of size can be roughly explained by the depth/size trade-off results known in deep learning theory. As a matter of fact, many theoretical research papers have underlined the benefits of depth relatively to width for parameterizing complex functions (as, for example, in Telgarsky, 2015, 2016). For a fixed the number of neurons, when comparing two neural networks, the deepest is the most expressive one (Lu et al., 2017).

It turns out that Theorem 1 can be significantly refined when the partition W satisfies some geometrical properties. Our next proposition examines the case where the sets W_1, \dots, W_{m_f} are convex.

Corollary 1. Let $f \in \text{Lip}_1(\mathbb{R}^d)$ that is also m_f -piecewise linear with convex subdomains W_1, \dots, W_{m_f} . Then there exists a neural network of the form (1) verifying Assumption 1 that represents f . Besides, its depth is $d \log_2(m_f) + 1$ and its size is at most $3m_f^2 + m_f - 1$.

Corollary 1 offers a significant improvement over Theorem 1, since in general $M_f = m_f$. We note in passing that the result of this proposition is dimension-free.

3.2 GroupSort neural networks on the real line

Piecewise linear functions defined on a convex set deserve a special treatment, since in this case, any connected subset is convex.

Proposition 2. Let $f \in \text{Lip}_1(\mathbb{R})$ that is also m_f -piecewise linear. Then there exists a neural network of the form (1)

verifying Assumption 1 that represents f . Besides, its depth is $2d \log_2(m_f) + 1$ and its size is at most $3m_f^2 + m_f - 1$.

In the specific case where f is convex (or concave), then there exists a neural network of the form (1) verifying Assumption 1 that represents f . Its depth is $d \log_2(m_f) + 1$ and its size is at most $3m_f - 1$.

When f is convex (or concave) and $d = 2^n$ for some $n > 1$, then there exists a neural network of the form (1) verifying Assumption 1 that represents f . Its depth is $d \log_2(m_f) + 1$ and its size is at most $3m_f - 1$.

This proposition is the counterpart of Arora et al. (2018, Theorem 2.2), which states that any m_f -piecewise linear function from $\mathbb{R} \rightarrow \mathbb{R}$ can be represented by a 2-layer ReLU neural network with a size at least $m_f - 1$. He et al. (2018, Theorem 5.2) shows that the upper-bound on the size of ReLU networks is $O(2^{m^2 + 2(m-1)})$. Thus, for the representation of piecewise linear functions on the real line, GroupSort networks require larger depths but smaller sizes. Besides, bear in mind that the obtained ReLU neural networks do not necessarily verify a requirement similar to the one of Assumption 1.

Regarding the number of linear regions of GroupSort networks on the real line, we have the following result:

Lemma 3. Any neural network of the form (1) on the real line, with depth q and widths n_1, \dots, n_{q-1} , parameterizes a piecewise linear function with at most $2^{q-2} \prod_{i=1}^{q-1} n_i$ linear subdomains.

We deduce from this lemma that for a neural network of the form (1) with depth $q > 2$ and constant width m , the maximum number of linear regions is $O(2^{q-3} m^{q-1})$. Similarly to ReLU networks (Montúfar et al., 2014; Arora et al., 2018), the maximum number of linear regions for GroupSort networks with grouping size 2 is also likely to grow polynomially in n and exponentially in m .

Our next corollary now illustrates the trade-off between depth and width for GroupSort neural networks.

Corollary 2. Let $f \in \text{Lip}_1(\mathbb{R})$ be an m_f -piecewise linear function. Then, any neural network of the form (1) verifying Assumption 1 and representing f with a depth q , has a size at least $\frac{1}{2}(q-1)m_f^{1-(q-1)}$.

The lower bound highlighted in Corollary 2 is dependent on the depth q of the neural network. By looking at the minimum of the function, we get that any neural network representing f has a size $s > \frac{e \ln(m_f)}{2}$. Thus, merging this result with Proposition 2, we have that for any m_f -piecewise linear function from $\mathbb{R} \rightarrow \mathbb{R}$, there exists a GroupSort network verifying Assumption 1 with a size satisfying

$$\frac{e \ln(m_f)}{2} - 6 \leq 3m_f^2 - m_f - 3:$$

We realize that this inequality is large but, up to our knowledge, this is the first of this type for GroupSort neural networks.

3.3 Approximating Lipschitz continuous functions on compact sets

Following our plan, we tackle in this subsection the task of approximating Lipschitz continuous functions on compact sets using GroupSort neural networks. The space of continuous functions on $[0; 1]^d$ is equipped with the uniform norm

$$\|f - g\|_\infty = \max_{x \in [0; 1]^d} |f(x) - g(x)|$$

The main result of the section, and actually of the article, is that GroupSort neural networks are well suited for approximating functions in $\text{Lip}([0; 1]^d)$.

Theorem 2. Let $\epsilon > 0$ and $d > 2$, $f \in \text{Lip}_1([0; 1]^d)$. Then there exists a neural network of the form (1) verifying Assumption 1 such that $\|f - D\|_\infty \leq 6\epsilon$. The depth of D is $O(d^2 \log_2(\frac{2}{\epsilon}))$ and its size is $O((\frac{2}{\epsilon})^{d^2})$.

To the best of our knowledge, Theorem 2 is the first one that provides an upper bound on the depth and size of neural networks, with constraints on the weights, that approximate Lipschitz continuous functions.

As for the representation of piecewise linear functions, one can, for the sake of completeness, compare this bound with those previously found in the literature of ReLU neural networks. Yarotsky (2017) establishes the density of ReLU networks in Sobolev spaces, using a different technique of proof. In particular, Theorem 1 of this paper states that for any $f \in \text{Lip}_1([0; 1]^d)$ continuously differentiable, there exists a ReLU neural network approximating f with precision ϵ , with depth at most $(\ln(1/\epsilon) + 1)$ and size at most $c\epsilon^{-d(\ln(1/\epsilon) + 1)}$ (with a constant function of d). Comparing this result with our Theorem 2, we see that, with respect to ϵ , both depths are similar but ReLU networks are smaller in size. However, one has to keep in mind that both lines of proof largely differ. Besides, our formulation ensures that the approximator is also Lipschitz function, a feature that cannot be guaranteed under the formulation of Yarotsky (2017).

It turns out however that our framework provides smaller neural networks as soon as $d \geq 1$.

Proposition 3. Let $\epsilon > 0$ and $f \in \text{Lip}_1([0; 1])$. Then there exists a neural network D of the form (1) verifying Assumption 1 such that $\|f - D\|_\infty \leq 6\epsilon$. The depth of D is $2d \log_2(1/\epsilon) + 1$ and its size is $O((\frac{1}{\epsilon})^2)$.

Besides, if f is assumed to be convex or concave, then the depth of D is $d \log_2(1/\epsilon) + 1$ and its size is $O(\frac{1}{\epsilon})$.

4 Impact of the grouping size

The previous section paved the way for a better understanding of GroupSort neural networks and their ability to approximate Lipschitz continuous functions. As mentioned in Section 2, one can play with the grouping size of the neural network when defining its architecture. However, it is not clear how changing this parameter might influence the expressivity of the network. The present section aims at bringing some understanding. Following a similar reasoning as in Section 3, we start by analyzing how GroupSort networks with an arbitrary grouping size $k \geq 2$ can represent any piecewise linear functions:

Proposition 4 (Extension of Proposition 1) Let $f_1, \dots, f_m : \mathbb{R}^d \rightarrow \mathbb{R}$ be a collection of functions ($m > 2$), each represented by a neural network of the form (1), with common depth q and sizes $s_i = 1, \dots, m$.

In the specific case where $m = k^n$ for some $n > 1$, there exist neural networks of the form (1) (with grouping size k) with depth $q + \log_k(m)$ and size at most $s_1 + \dots + s_m + \frac{m-1}{k-1}$ that represent the functions $f = \max(f_1, \dots, f_m)$ and $g = \min(f_1, \dots, f_m)$.

Similarly to Section 3, this leads to the following corollary:

Corollary 3 (Extension of Corollary 1) Let $f \in \text{Lip}_1(\mathbb{R}^d)$ that is also m_f -piecewise linear with convex subdomains W_1, \dots, W_{m_f} such that $m_f = k^n$ for some $n > 1$. Then there exists a neural network of the form (1) verifying Assumption 1 that represents f . Besides, its depth is at most $d \log_k(m_f) + 1$ and its size is at most $\frac{m_f^2}{k-1}$.

Proposition 4 and Corollary 3 exhibit the nice properties of using larger grouping sizes. Indeed, for a given n , there exists a neural network with depth $2n + 1$ and grouping size k representing a function with k^n pieces. Consequently, the use of larger grouping sizes helps have more expressive neural networks. The efficiency of larger grouping sizes may also be explained by the following result for GroupSort networks on the real line:

Lemma 4 (Extension of Lemma 3) Any neural network of the form (1) on the real line, with depth q , widths n_1, \dots, n_{q-1} , and grouping size k , parameterizes a piecewise linear function with at most $k^{q-2} \cdot (\frac{(k-1)n_1}{2} + 1)$ linear subdomains.

Thus, the number of linear regions of a GroupSort network is likely to increase polynomially with the grouping size, which highlights the benefits of using larger groups. Similarly to Section 3, when moving to the approximation of Lipschitz continuous functions on $[0; 1]^d$, we are led to the following theorem:

Theorem 3 (Extension Theorem 2) Let $\epsilon > 0$, $d > 2$, and $f \in \text{Lip}_1([0; 1]^d)$. Then there exists a neural network of the form (1) verifying Assumption 1 with grouping size

Methods	Up Depth	Up Size	Down Size	Reference
Representing $m = k^n$ -PWL functions in \mathbb{R}^d with a constant width n				
ReLU	$d \log_2(d+1)e + 1$	$O(d2^{m^2})$	$O(m)$	He et al. (2018)
GroupSortGS = k	$d \log_k(m)e + 1$	$\frac{m^2 - 1}{k - 1}$	$\frac{n \log_k(m)}{2 \log_k(n)}$	present article
Approximating 1-Lipschitz continuous functions in $[0, 1]^d$				
ReLU	$O(\ln(\frac{1}{\epsilon}))$	$O(\frac{\ln(1-\epsilon)}{\epsilon^d})$	n	Yarotsky (2017)
GroupSortGS = $d \frac{2^p - d}{e} e$	$O(d^2)$	$O((\frac{2^p - d}{e})^d - 1)$	n	present article
Approximating 1-Lipschitz continuous functions in $[0, 1]$				
ReLU (PWL representation)	2	$O(2^{1-\epsilon^2+2-\epsilon})$	n	He et al. (2018)
ReLU (different approach)	$O(\ln(\frac{1}{\epsilon}))$	$O(\frac{\ln(1-\epsilon)}{\epsilon})$	n	Yarotsky (2017)
Adaptative ReLU	6	$O(\frac{1}{e \ln(1-\epsilon)})$	n	Yarotsky (2017)
GroupSortGS = $d \frac{1}{e} e$	3	$O(\frac{1}{e})$	n	present article

Table 1: Summary of the results shown in the present paper together with results previously found for ReLU networks. “Up Depth” refers to upper bounds on the depths, “Up Size” to upper bounds on the sizes, and “Down Size” to lower bounds on the sizes. The symbol “?” means that no result is known (up to our knowledge).

$d \frac{2^p - d}{e} e$ such that f is D -Lipschitz. The depth of D is $O(d^2)$ and its size is $O((\frac{2^p - d}{e})^d - 1)$.

Using a grouping size proportional to $\frac{1}{\epsilon}$, we thus have a bound on the depth that is independent from the error rate. The uni-dimensional case leads to a different result:

Proposition 5 (Extension of Proposition 3) Let $\epsilon > 0$ and $f \in \text{Lip}_1([0, 1])$. Then there exists a neural network of the form (1) verifying Assumption 1 (with grouping size k) such that f is D -Lipschitz. The depth of D is $2 \log_k(\frac{1}{\epsilon})e + 1$ and its size is at most $O(\frac{1}{k\epsilon^2})$.

In particular, if k is chosen to be equal to $\frac{1}{\epsilon}$, then the depth of D is 3 and its size is $O(\frac{1}{\epsilon})$.

When approximating real-valued functions, the use of large grouping sizes can significantly decrease the required size since it goes from $O(1-\epsilon^2)$ in Proposition 3 to $O(1-\epsilon)$ in Proposition 5. When f is assumed to be convex or concave, the depth of the network can further be reduced to 2.

Using a different approach for approximating Lipschitz continuous functions in $[0, 1]$, Yarotsky (2017, Theorem 1) shows that ReLU networks with a depth $O(\ln(1-\epsilon))$ is needed together with a size $O(\frac{\ln(1-\epsilon)}{\epsilon})$ to approximate with an error rate ϵ . To sum-up, when compared with ReLU networks, GroupSort neural networks with well-chosen grouping size can be significantly more expressive.

Table 1 summarizes the results shown in the present paper together with results previously found for ReLU networks. Bear in mind that GroupSort neural networks also have the supplementary condition that any parameterized function verifies the 1-Lipschitz continuity.

5 Experiments

Anil et al. (2019) have already compared the performances of GroupSort neural networks with their ReLU counterparts, both with constraints on the weights. In particular, they showed that ReLU neural networks are more sensitive to adversarial attacks while stressing the fact that if their weights are limited, then these networks lose their expressive power. Building on these observations, we further illustrate the good behavior of GroupSort neural networks in the context of estimating a Lipschitz continuous regression function and in approximating the Wasserstein distance (via its dual form) between pairs of distributions.

Impact of the depth. We start with the problem of learning a function f in the model $Y = f(X)$, where X follows a uniform distribution on $[-\delta, \delta]$ and f is 32-piecewise linear. For this aim, we use neural networks of the form (1) with respective depths $q = 2, 8, 14, 20$, and a constant width $n = 50$. Since we are only interested in the approximation properties of the networks, we assume to have at hand an infinite number of pairs $(X_i; f(X_i))$ and train the models by minimizing the mean squared error. We give in the Appendix, the full details of our experimental setting. The quality of the estimation is evaluated using the uniform norm between the target function f and the output network. In order to enforce Assumption 1, GroupSort neural networks are constrained using the orthonormalization of Björck and Bowie (1971). The results are presented in Figure 3. Note that throughout this section, confidence intervals are computed over 20 runs. In line with Theorem 1, which states that f is representable by a neural network of the form (1) with size at most $3 \cdot 32^2 + 32 - 1 = 3104$, we clearly observe that, as the depth of the networks increases, the uniform norm de-




Figure 3: Reconstruction of a 32-piecewise linear function on $[-8; 8]$ with a GroupSort neural network of the form (1) with depth $q = 2, 8, 14, 20$, and a constant width $n = 50$ (the thickness of the line represents a 95%-confidence interval).

creases and the Lipschitz constant of the network converges to 1. The reconstruction of this piecewise linear function is even almost perfect for the depth $q = 20$, i.e., with a network of size only $20 \cdot 60 = 1200$, a value significantly smaller than the upper bound of the theorem.

We also illustrate the behavior of GroupSort neural networks in the context of WGANs (Arjovsky et al., 2017). We run a series of small experiments in the simplified setting where we try to approximate the 1-Wasserstein distance between two bivariate mixtures of independent Gaussian distributions with 4 components. We consider networks of the form (1) with grouping size 2, a depth $q = 2$ and $q = 5$, and a constant width $n = 20$. For a pair of distributions $(m; n)$, our goal is to exemplify the relationship between the 1-Wasserstein distance $\sup_{f \in \mathcal{L}_{\text{Lip}_1}(\mathbb{R}^2)} (E_m - E_n)$ (approximated with the Python package by Flamary and Courty, 2017) and the neural distance $\sup_{f \in \mathcal{D}_2} (E_m - E_n)$ (Arora et al., 2017) computed over the class of functions \mathcal{D}_2 . To this aim, we randomly draw 40 different pairs of distributions. Then, for each of these pairs, we compute an approximation of the 1-Wasserstein distance and calculate the corresponding neural distance. Figure 4 depicts the best parabolic fit between 1-Wasserstein and neural distances, and shows the corresponding Least Relative Error (LRE) together with the width of the envelope. The take-home message of this figure is that both the LRE and the width are significantly smaller for deeper GroupSort neural networks.

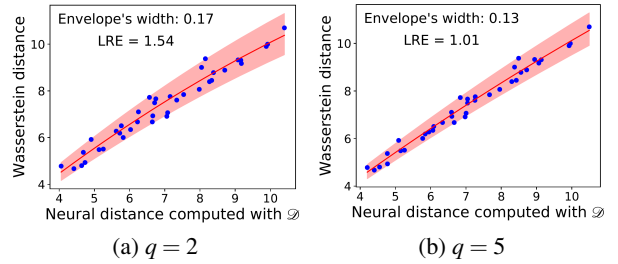


Figure 4: Scatter plots of 40 pairs of Wasserstein and neural distances computed with GroupSort neural networks, for $q = 2; 5$. The underlying distributions are bivariate Gaussians. The red curve is the optimal parabolic fitting and LRE refers to the Least Relative Error. The red zone is the envelope obtained by stretching the optimal curve.

Impact of the grouping size. To highlight the benefits of using larger grouping sizes, we show the impact of increasing the grouping size from 2 in Figure 5a to 5 in Figure 5b for the representation of a 20-piecewise linear function. This is corroborated by Figure 5c, which illustrates that the uniform norm with a 64-piecewise linear function decreases when the grouping size increases. As already underlined in Lemma 4, this may be explained by the fact that the number of linear regions significantly grows with the grouping size—see Figure 5d.

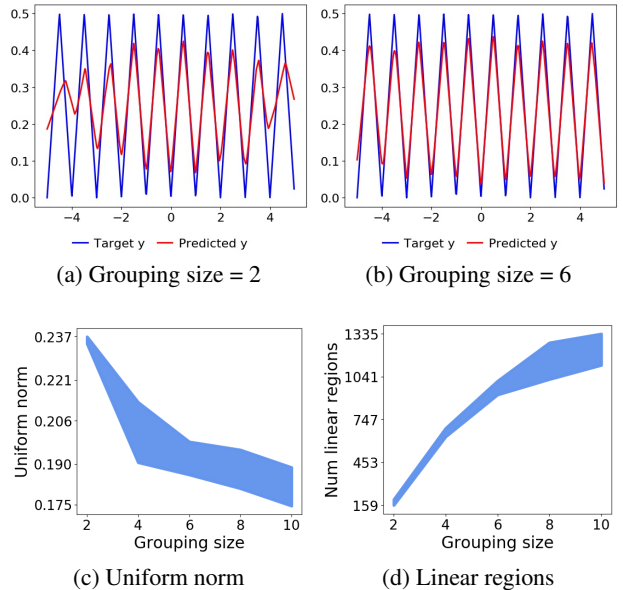


Figure 5: Reconstruction of a 20-piecewise linear function on $[-5; 5]$ (top line) and a 64-piecewise linear function (bottom line) with GroupSort neural networks of the form (1) with depth $q = 4$ and varying grouping sizes $k = 2; 4; 6; 8; 10$.

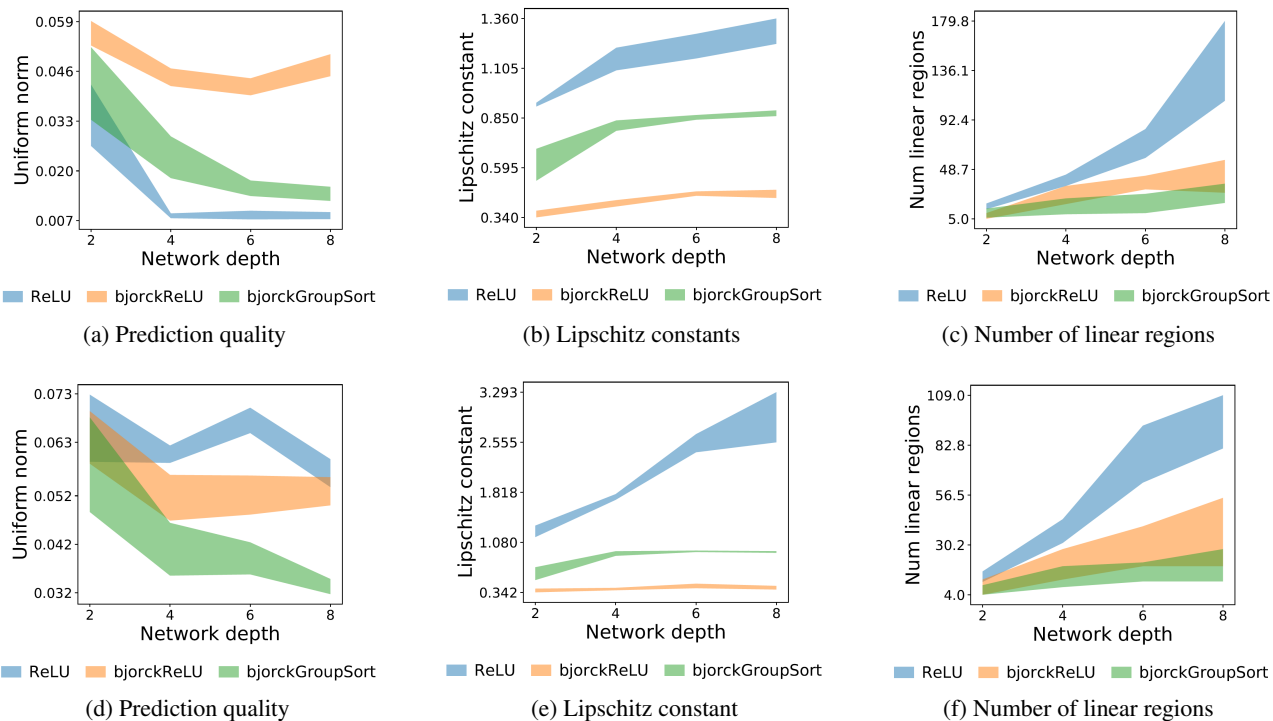


Figure 6: (Top line) Estimating the function $f(x) = (1=15) \sin(15x)$ on $[0;1]$ in the model $Y = f(X)$, with a dataset of size $n = 100$. (Bottom line) Estimating the function $f(x) = (1=15) \sin(15x)$ on $[0;1]$ in the model $Y = f(X) + e$, with a dataset of size $n = 100$ (the thickness of the line represents a 95%-confidence interval).

Comparison with ReLU neural networks. Next, in a second series of experiments, we compare the performances of GroupSort networks against two baselines: ReLU neural networks without constraints on the weights (dense in the set of continuous functions on a compact set; see Yarotsky, 2017), and ReLU neural networks with orthonormalization of Björck and Bowie (1971). The architecture of the ReLU neural networks in terms of depth and width is the same as for GroupSort networks: $q = 2, 4, 6, 8$, and $w = 20$. The task is now to approximate the 1-Lipschitz continuous function $f(x) = (1=15) \sin(15x)$ on $[0;1]$ in the models $Y = f(X)$ (noiseless case) and $Y = f(X) + e$ (noisy case), where X is uniformly distributed on $[0;1]$ and e follows a Gaussian distribution with standard deviation 0.05. In both cases, we assume to have at hand a finite sample of size $n = 100$ and fit the models by minimizing the mean squared error.

Both results (noiseless case and noisy case) are presented in Figure 6. We observe that in the noiseless setting Figure 6a, 6b, and 6c, ReLU neural networks without normalization have a slightly better performance with respect to the uniform norm with, however, a Lipschitz constant larger than 1. On the other hand, in the noisy case, ReLU neural networks without constraints have a tendency to overfitting (a high Lipschitz constant close to 2.7), leading to a deteriorated performance, contrary to GroupSort neural networks. Furthermore, in both cases (noiseless and noisy), ReLU with

constraints are found to perform worse (due to a Lipschitz constant much smaller than 1) than their GroupSort counterparts in terms of prediction. Interestingly, we see in the two examples shown in Figure 6e and Figure 6f, that the number of linear regions for GroupSort neural networks is smaller than for ReLU networks.

Finally, we quickly show in Appendix a comparison between GroupSort and ReLU networks when approximating Wasserstein distances. The take home message is that, on this specific task, GroupSort networks perform better.

6 Conclusion

The results presented in this article show the advantage of using GroupSort neural networks over standard ReLU networks. On the one hand, ReLU neural networks without any constraints are sensitive to adversarial attacks (as they may have a large Lipschitz constant) and, on the other hand, lose expressive power when enforcing limits on their weights. On the opposite, GroupSort neural networks with constrained weights are proved to be both robust and expressive, and are therefore an interesting alternative. Moreover, by allowing larger grouping sizes for GroupSort networks, one can further increase their expressivity. These properties open new perspectives for broader use of GroupSort networks.

References

- Anil, C., Lucas, J., and Grosse, R. (2019). Sorting out Lipschitz function approximation. In Chaudhuri, K. and Salakhutdinov, R., editors, *Proceedings of the 36th International Conference on Machine Learning*, volume 97, pages 291–301. PMLR.
- Arjovsky, M., Chintala, S., and Bottou, L. (2017). Wasserstein generative adversarial networks. In Precup, D. and Teh, Y., editors, *Proceedings of the 34th International Conference on Machine Learning*, volume 70, pages 214–223. PMLR.
- Arora, R., Basu, A., Mianjy, P., and Mukherjee, A. (2018). Understanding deep neural networks with rectified linear units. In *International Conference on Learning Representations*.
- Arora, S., Ge, R., Liang, Y., Ma, T., and Zhang, Y. (2017). Generalization and equilibrium in generative adversarial nets (GANs). In *Proceedings of the 34th International Conference on Machine Learning-Volume 70*, pages 224–232.
- Biau, G., Cadre, B., Sangnier, M., and Tanielian, U. (2020). Some theoretical properties of GANs. *The Annals of Statistics*, 48:1539–1566.
- Björck, A. and Bowie, C. (1971). An iterative algorithm for computing the best estimate of an orthogonal matrix. *SIAM Journal on Numerical Analysis*, 8:358–364.
- Blanchet, J., Kang, Y., and Murthy, K. (2019). Robust Wasserstein profile inference and applications to machine learning. *Journal of Applied Probability*, 56:830–857.
- Cybenko, G. (1989). Approximation by superpositions of a sigmoidal function. *Mathematics of Control, Signals and Systems*, 2:303–314.
- Devroye, L., Györfi, L., and Lugosi, G. (1996). *A Probabilistic Theory of Pattern Recognition*. Springer, New York.
- Esfahani, P. and Kuhn, D. (2018). Data-driven distributionally robust optimization using the Wasserstein metric: Performance guarantees and tractable reformulations. *Mathematical Programming*, 171:115–166.
- Flamary, R. and Courty, N. (2017). POT: Python Optimal Transport library.
- Gao, R., Chen, X., and Kleywegt, A. (2017). Wasserstein distributional robustness and regularization in statistical learning. *arXiv:1712.06050*.
- Goodfellow, I., Shlens, J., and Szegedy, C. (2015). Explaining and harnessing adversarial examples. In *International Conference on Learning Representations*.
- Gulrajani, I., Ahmed, F., Arjovsky, M., Dumoulin, V., and Courville, A. (2017). Improved training of Wasserstein GANs. In Guyon, I., von Luxburg, U., Bengio, S., Wallach, H., Fergus, R., Vishwanathan, S., and Garnett, R., editors, *Advances in Neural Information Processing Systems 30*, pages 5767–5777. Curran Associates, Inc.
- He, J., Li, L., Xu, J., and Zheng, C. (2018). ReLU deep neural networks and linear finite elements. *arXiv:1807.03973*.
- Hornik, K., Stinchcombe, M., and White, H. (1989). Multilayer feedforward networks are universal approximators. *Neural Networks*, 2:359–366.
- Huster, T., Chiang, C.-Y. J., and Chadha, R. (2018). Limitations of the Lipschitz constant as a defense against adversarial examples. In *Joint European Conference on Machine Learning and Knowledge Discovery in Databases*, pages 16–29. Springer.
- Kodali, N., Abernethy, J., Hays, J., and Kira, Z. (2017). On convergence and stability of GANs. *arXiv:1705.07215*.
- Lu, Z., Pu, H., Wang, F., Hu, Z., and Wang, L. (2017). The expressive power of neural networks: A view from the width. In *Advances in Neural Information Processing Systems*, pages 6231–6239.
- Miyato, T., Kataoka, T., Koyama, M., and Yoshida, Y. (2018). Spectral normalization for generative adversarial networks. In *International Conference on Learning Representations*.
- Montúfar, G., Pascanu, R., Cho, K., and Bengio, Y. (2014). On the number of linear regions of deep neural networks. In Ghahramani, Z., Welling, M., Cortes, C., Lawrence, N., and Weinberger, K., editors, *Advances in Neural Information Processing Systems 27*, pages 2924–2932. Curran Associates, Inc.
- Raghu, M., Poole, B., Kleinberg, J., Ganguli, S., and Dickstein, J. (2017). On the expressive power of deep neural networks. In Precup, D. and Teh, Y., editors, *Proceedings of the 34th International Conference on Machine Learning*, volume 70, pages 2847–2854. PMLR.
- Telgarsky, M. (2015). Representation benefits of deep feedforward networks. *arXiv 1509.08101*.
- Telgarsky, M. (2016). Benefits of depth in neural networks. In Feldman, V., Rakhlin, A., and Shamir, O., editors, *29th Annual Conference on Learning Theory*, volume 49, pages 1517–1539. PMLR.
- Villani, C. (2008). *Optimal Transport: Old and New*. Springer, Berlin.
- Wei, X., Gong, B., Liu, Z., Lu, W., and Wang, L. (2018). Improving the improved training of Wasserstein GANs: A consistency term and its dual effect. *arXiv:1803.01541*.
- Yarotsky, D. (2017). Error bounds for approximations with deep ReLU networks. *Neural Networks*, 94:103–114.
- Zhou, Z., Liang, J., Song, Y., Yu, L., Wang, H., Zhang, W., Yu, Y., and Zhang, Z. (2019). Lipschitz generative adversarial nets. In Chaudhuri, K. and Salakhutdinov,

R., editors, *Proceedings of the 36th International Conference on Machine Learning*, volume 97, pages 7584–7593. PMLR.