

On the CVP for the root lattices via folding with deep ReLU neural networks

Vincent Corlay^{†,*}, Joseph J. Boutros[‡], Philippe Ciblat[†], and Loïc Brunel^{*}

[†] Telecom ParisTech, 46 Rue Barrault, 75013 Paris, France, v.corlay@fr.mercede.mee.com

[‡] Texas A&M University, Doha, Qatar, *Mitsubishi Electric R&D Centre Europe, Rennes, France

Abstract—Point lattices and their decoding via neural networks are considered in this paper. Lattice decoding in \mathbb{R}^n , known as the closest vector problem (CVP), becomes a classification problem in the fundamental parallelotope with a piecewise linear function defining the boundary. Theoretical results are obtained by studying root lattices. We show how the number of pieces in the boundary function reduces dramatically with folding, from exponential to linear. This translates into a two-layer ReLU network requiring a number of neurons growing exponentially in n to solve the CVP, whereas this complexity becomes polynomial in n for a deep ReLU network.

I. INTRODUCTION AND MOTIVATIONS

The objective of this paper is two-fold. Firstly, we introduce a new paradigm to solve the CVP. This approach enables to find efficient decoding algorithms for some dense lattices. For instance, such a neural network for the Gosset lattice is a key component of a neural Leech decoder. Secondly, we also aim at contributing to the understanding of the efficiency of deep learning, namely the expressive power of deep neural networks. As a result, our goal is to present new decoding algorithms and interesting functions that can be efficiently computed by deep neural networks.

Deep Learning is about two key aspects: (i) Finding a function class $\Phi = \{f\}$ that contains a function f^* “close enough” to a target function. (ii) Finding a learning algorithm L for the class Φ . Of course the choices of (i) and (ii) can be either done jointly or separately but in either case they impact each other. Research on the expressive power of deep neural networks focuses mostly on (i) [7][11], by studying some specific functions contained in the function class of a network. Typically, the aim is to show that there exist functions that can be well approximated by a deep network with a polynomial number of parameters whereas an exponential number of parameters is required for a shallow network. This line of work leads to “gap” theorems and “capacity” theorems for deep networks and it is similar to the classical theory of Boolean circuit complexity [8].

In this scope, several papers investigate specifically deep ReLU networks [10][9][11][14][12][2] (See [9, Section 2.1] for a short introduction to ReLU neural networks). Since a ReLU network computes a composition of piecewise affine functions, all functions in Φ are continuous piecewise linear (CPWL). Hence, the efficiency of Φ can be evaluated by checking whether a CPWL function with a lot of affine pieces belongs to Φ . For example, there exists at least one function in \mathbb{R}^n with $\Omega\left(\left(w/n\right)^{L-1}w^n\right)$ affine pieces that can be computed

with a w -wide deep ReLU network having L hidden layers [9]. A two-layer network would need an exponential number of parameters for this same function. In [11] they further show that any random deep network achieves a similar exponential behavior.

Some results in the literature are established by considering elementary oscillatory functions (see e.g. [15] - VII-F) or piecewise linear functions with regions of random shape. It is not clear whether such types of functions may arise naturally in computer science and engineering fields.

Our work lies somewhere between [9][14] and [11]: our functions are neither elementary nor random. We discovered them in the context of sphere packing and lattices which are solution to many fundamental problems in number theory, chemistry, communication theory, string theory, and cryptography [4]. Hence, in contrary to existing works, we do not search for a specific class of functions to justify the success of deep learning. We set technical problems, study the functions arising from lattices, and show that deep networks are suited to tackle them as we obtain similar gap theorems between shallow and deep models.

II. MAIN RESULTS

A ReLU network with a finite number of neurons is not capable to infer infinite periodic functions (see [15] - VII-F). Hence, it cannot implement a simple modulo operation. As a result, we allow a low complexity pre-processing of the point to be decoded to obtain an equivalent position in the fundamental parallelotope $\mathcal{P}(\mathcal{B})$ of the lattice.

- 1) Theorems 1&2 show that the decision boundary for the hyperplane logical decoder (HLD) [5] is given by a continuous piecewise linear function, for any lattice with a Voronoi-reduced basis. Corollary 1 guarantees the same result for a semi-Voronoi-reduced basis.
- 2) For the lattice A_n with a basis defined by the Gram matrix (4) and a point in $\mathcal{P}(\mathcal{B})$, Theorem 4 proves that the decision function has $\Omega(2^n)$ affine pieces.
- 3) Also for A_n , the number of pieces is reduced to $\mathcal{O}(n)$ after folding as stated in Theorem 5.
- 4) Results of Section V-B and V-C, based on Theorem 5, implies that there exists a ReLU network of depth $\mathcal{O}(n^2)$ and width $\mathcal{O}(n^2)$ solving the CVP.
- 5) Theorem 6 shows that a ReLU network with only one hidden layer needs $\Omega(2^n)$ neurons to solve the CVP.

Moreover, the theory presented in this paper is not limited to A_n . It extends very well to other dense lattices. Indeed, we

already obtained similar results for all root lattices. They will however be presented in future communications due to lack of space. Finally, this paradigm seems not to be limited to lattices as it may extend to binary block codes (see e.g. Figure 4 in [15] - VII-F).

III. LATTICES AND POLYTOPES

A lattice Λ is a discrete additive subgroup of \mathbb{R}^n . For a rank- n lattice in \mathbb{R}^n , the rows of a $n \times n$ generator matrix G constitute a basis of Λ and any lattice point x is obtained via $x = zG$, where $z \in \mathbb{Z}^n$. For a given basis $\mathcal{B} = \{b_i\}_{i=1}^n$, $\mathcal{P}(\mathcal{B})$ denotes the fundamental parallelotope of Λ and $\mathcal{V}(x)$ the Voronoi cell of a lattice point x [5]. The minimum Euclidean distance of Λ is $d_{\min}(\Lambda) = 2\rho$, where ρ is the packing radius.

A vector $v \in \Lambda$ is called Voronoi vector if the half-space $\{y \in \mathbb{R}^n : y \cdot v \leq \frac{1}{2}v \cdot v\}$ has a non empty intersection with $\mathcal{V}(0)$. The vector is said relevant if the intersection is an $n-1$ -dimensional face of $\mathcal{V}(0)$. We denote by τ_f the number of relevant Voronoi vectors, referred to as the *Voronoi number* in the sequel. For root lattices [4], the Voronoi number is equal to the kissing number τ . For random lattices, we typically have $\tau_f = 2^{n+1} - 2$. The set $\tau_f(x)$, for $x \in \Lambda$, is the set of lattice points having a common Voronoi facet with x .

The next definition, introduced in [5], is important for the rest of the paper.

Definition 1. Let \mathcal{B} be the \mathbb{Z} -basis of a rank- n lattice Λ in \mathbb{R}^n . \mathcal{B} is said *Voronoi-reduced (VR)* if, for any point $y \in \mathcal{P}(\mathcal{B})$, the closest lattice point \hat{x} to y is one of the 2^n corners of $\mathcal{P}(\mathcal{B})$, i.e. $\hat{x} = \hat{z}G$ where $\hat{z} \in \{0, 1\}^n$.

Lattice decoding refers to the method of finding the closest lattice point, the closest in Euclidean distance sense. This problem is also known as the closest vector problem. The neural lattice decoder employs $\mathcal{P}(\mathcal{B})$ as its main compact region [5], thus it is important to characterize $\mathcal{P}(\mathcal{B})$ as made below.

A k -dimensional element of $\mathcal{P}(\mathcal{B})$ is referred to as k -face. $\mathcal{P}(\mathcal{B})$ has 2^n 0-faces, called corners or vertices. This set of corners is denoted $\mathcal{C}_{\mathcal{P}(\mathcal{B})}$. Moreover, the subset of $\mathcal{C}_{\mathcal{P}(\mathcal{B})}$ obtained with $z_i = 1$ is $\mathcal{C}_{\mathcal{P}(\mathcal{B})}^1$ and $\mathcal{C}_{\mathcal{P}(\mathcal{B})}^0$ for $z_i = 0$. The remaining faces of $\mathcal{P}(\mathcal{B})$ are parallelotopes. For instance, a $n-1$ -dimensional facet of $\mathcal{P}(\mathcal{B})$, say \mathcal{F}_i , is itself a parallelotope of dimension $n-1$ defined by $n-1$ vectors of \mathcal{B} . Throughout the paper, the term facet refers to a $n-1$ -face.

A polyhedron is defined as the intersection of a finite number of half-spaces generated by hyperplanes:

$$P_o = \{x \in \mathbb{R}^n : xA \leq b, A \in \mathbb{R}^{n \times m}, b \in \mathbb{R}^m\}.$$

A polyhedron is *convex* if none of the bounding hyperplanes crosses its interior. A convex polyhedron is called a *polytope* [6]. In this paper, we use not only parallelotopes but also simplices. A n -simplex associated to \mathcal{B} is given by

$$S(\mathcal{B}) = \{y \in \mathbb{R}^n : y = \sum_{i=0}^n \alpha_i b_i, \sum_{i=1}^n \alpha_i \leq 1, \alpha_i \geq 0 \forall i\}.$$

It is clear that the corners of $S(\mathcal{B})$, the set $\mathcal{C}_{S(\mathcal{B})}$, are the $n+1$ points $\{0, b_1, \dots, b_n\}$.

We say that a function $g : \mathbb{R}^{n-1} \rightarrow \mathbb{R}$ is continuous piecewise linear (CPWL) if there exists a finite set of polyhedra covering \mathbb{R}^{n-1} (which implies continuity), and g is affine over each polyhedron. The number of pieces of g is the number of distinct expressions of the local affine functions.

Finally, \vee and \wedge denote respectively the maximum and the minimum operator. We define a convex (resp. concave) CPWL function formed by a set of affine functions related by the operator \vee (resp. \wedge). If $\{g_k\}$ is a set of K affine functions, the function $f = g_1 \vee \dots \vee g_K$ is CPWL and convex.

IV. THE DECISION BOUNDARY FUNCTION

Given a VR basis, after translating the point to be decoded inside $\mathcal{P}(\mathcal{B})$, the HLD decoder proceeds in estimating each z_i -component separately [5]. The HLD computes the position of y relative to a boundary via a Boolean equation to guess whether $z_i = 0$, i.e. the closest lattice point belongs to $\mathcal{C}_{\mathcal{P}(\mathcal{B})}^0$, or $z_i = 1$ when the closest lattice point is in $\mathcal{C}_{\mathcal{P}(\mathcal{B})}^1$. This boundary cuts $\mathcal{P}(\mathcal{B})$ into two regions. It is composed of Voronoi facets of the corner points. The next step is to study the decision boundary function. **For the rest of the paper, without loss of generality, the integer coordinate to be decoded is z_1 .**

We recall that a variable $u_j(y)$ in the Boolean equations of the HLD is obtained as:

$$u_j(y) = \text{sign}(y \cdot v_j - p_j) \in \{0, 1\}, \quad (1)$$

where v_j is the orthogonal vector to the boundary hyperplane $\{y \in \mathbb{R}^n : y \cdot v_j - p_j = 0\}$. The latter contains the Voronoi facet of a point $x \in \mathcal{C}_{\mathcal{P}(\mathcal{B})}^1$ and a point from $\tau_f(x) \cap \mathcal{C}_{\mathcal{P}(\mathcal{B})}^0$. The decision boundary cutting $\mathcal{P}(\mathcal{B})$ into two regions, with $\mathcal{C}_{\mathcal{P}(\mathcal{B})}^0$ on one side and $\mathcal{C}_{\mathcal{P}(\mathcal{B})}^1$ on the other side, is the union of these Voronoi facets. Each facet can be defined by an affine function over a compact subset of \mathbb{R}^{n-1} , and the decision boundary is locally described by one of these functions.

Let $\{e_i\}_{i=1}^n$ be the canonical orthonormal basis of the vector space \mathbb{R}^n . For $y \in \mathbb{R}^n$, the j -th coordinate is $y^j = y \cdot e_j$. Denote $\tilde{y} = (y_2, \dots, y_n) \in \mathbb{R}^{n-1}$ and let $\mathcal{H} = \{h_j\}$ be the set of affine functions involved in the decision boundary. The affine boundary function $h_j : \mathbb{R}^{n-1} \rightarrow \mathbb{R}$ is

$$h_j(\tilde{y}) = y^1 = \left(p_j - \sum_{k \neq 1} y_k v_j^k \right) / v_j^1. \quad (2)$$

For the sake of simplicity, in the sequel h_j shall denote the function defined in (2) or its associated hyperplane $\{y \in \mathbb{R}^n : y \cdot v_j - p_j = 0\}$ depending on the context.

Theorem 1. Consider a lattice defined by a VR basis $\mathcal{B} = \{b_i\}_{i=1}^n$. Suppose that the $n-1$ points $\mathcal{B} \setminus \{b_1\}$ belong to the hyperplane $\{y \in \mathbb{R}^n : y \cdot e_1 = 0\}$. Then, the decision boundary is given by a CPWL function $f : \mathbb{R}^{n-1} \rightarrow \mathbb{R}$, expressed as

$$f(\tilde{y}) = \wedge_{m=1}^M \{ \vee_{k=1}^{l_m} g_{m,k}(\tilde{y}) \}, \quad (3)$$

where $g_{m,k} \in \mathcal{H}$, $1 \leq l_m < \tau_f$, and $1 \leq M \leq 2^{n-1}$.




Fig. 1. CPWL decision boundary function for A_3 . The basis vectors are represented by the blue lines. The corner points in $\mathcal{C}_{\mathcal{P}(\mathcal{B})}^1$ are in red and the corner points in $\mathcal{C}_{\mathcal{P}(\mathcal{B})}^0$ in black.

In the next theorem, the orientation of the axes relative to \mathcal{B} does not require $\{b_i\}_{i=2}^n$ to be orthogonal to e_1 , neither b_1 and e_1 to be collinear.

Theorem 2. Consider a lattice defined by a VR basis $\mathcal{B} = \{b_i\}_{i=1}^n$. Without loss of generality, assume that $b_1^1 > 0$. Suppose also that $x^1 > \lambda^1, \forall x \in \mathcal{C}_{\mathcal{P}(\mathcal{B})}^1$ and $\forall \lambda \in \tau_f(x) \cap \mathcal{C}_{\mathcal{P}(\mathcal{B})}^0$. Then, the decision boundary is given by a CPWL function as in (3).

See [15] - VII-B for the proofs. Some interesting lattices may not admit a VR basis, e.g. see E_6 in [5]. In this case, if $\text{Vol}(\mathcal{P}(\mathcal{B}) \setminus \cup_{x \in \mathcal{C}(\mathcal{B})} \mathcal{V}(x)) \ll \text{Vol}(\mathcal{P}(\mathcal{B}))$ then HLD yields efficient decoding. A basis satisfying this condition is called *quasi-Voronoi-reduced*. The new definition below presumes that \mathcal{B} is quasi-Voronoi-reduced in order to make a successful discrimination of z_1 via the boundary function. Also, a surface in \mathbb{R}^n defined by a function g of $n-1$ arguments is written as $\text{Surf}(g) = \{(g(\tilde{y}), \tilde{y}) \in \mathbb{R}^n : \tilde{y} \in \mathbb{R}^{n-1}\}$.

Definition 2. Let \mathcal{B} be a basis of Λ . Assume that \mathcal{B} and $\{e_i\}_{i=1}^n$ have the same orientation as in Theorem 1. The basis is called *semi-Voronoi-reduced (SVR)* if there exists at least two points $x_1, x_2 \in \mathcal{C}_{\mathcal{P}(\mathcal{B})}^1$ such that $\text{Surf}(\bigvee_{k=1}^{\ell_1} g_{1,k}) \cap \text{Surf}(\bigvee_{k=1}^{\ell_2} g_{2,k}) \neq \emptyset$, where $\ell_1, \ell_2 \geq 1$, $g_{1,k}$ are the facets between x_1 and all points in $\tau_f(x_1) \cap \mathcal{C}_{\mathcal{P}(\mathcal{B})}^0$, and $g_{2,k}$ are the facets between x_2 and all points in $\tau_f(x_2) \cap \mathcal{C}_{\mathcal{P}(\mathcal{B})}^0$.

The above definition of a SVR basis imposes that the boundaries around two points of $\mathcal{C}_{\mathcal{P}(\mathcal{B})}^1$, defined by the two convex functions $\bigvee_{k=1}^{\ell_m} g_{m,k}$, $m = 1, 2$, have a non-empty intersection. Consequently, the min operator \wedge leads to a boundary function as in (3).

Corollary 1. $\mathcal{P}(\mathcal{B})$ for a SVR basis \mathcal{B} admits a decision boundary defined by a CPWL function as in (3).

Example 1. Consider the lattice A_3 defined by the Gram matrix (4). To better illustrate the symmetries we rotate the basis to have b_1 collinear with e_1 . Theorem 2 ensures that the decision boundary is a function. The function is illustrated on Figure 1 and its equation is (we omit the \tilde{y} in the formula to




Fig. 2. “Neighbor” figure of $\mathcal{C}_{\mathcal{P}(\mathcal{B})}$ for A_3 . Each edge connects a point $x \in \mathcal{C}_{\mathcal{P}(\mathcal{B})}^1$ to an element of $\tau_f(x) \cap \mathcal{C}_{\mathcal{P}(\mathcal{B})}^0$. The i edges connected to a point $x \in \mathcal{C}_{\mathcal{P}(\mathcal{B})}^1$ are the 1-faces of a regular i -simplex.

lighten the notations):

$$f = \left[h_{p_1} \vee h_1 \vee h_2 \right] \wedge \left[(h_{p_2} \vee h_1) \wedge (h_{p_2} \vee h_2) \right] \wedge \left[h_{p_3} \right],$$

where h_{p_1} , h_{p_2} and h_{p_3} are hyperplanes orthogonal to b_1 (the p index stands for plateau). On Figure 2 each edge is orthogonal to a local affine function of f and labeled accordingly. The $[\cdot]$ groups all the set of convex pieces of f that includes the same h_{p_j} . Functions for higher dimensions are available in [15] - VII-A.

From now on, **the default orientation of the basis with respect to the canonical axes of \mathbb{R}^n is assumed to be the one of Theorem 1.** We call f the decision boundary function. The domain of f (its input space) is $\mathcal{D} \subset \mathbb{R}^{n-1}$. The domain \mathcal{D} is the projection of $\mathcal{P}(\mathcal{B})$ on the hyperplane $\{e_i\}_{i=2}^n$. It is a bounded polyhedron that can be partitioned into disjoint convex (and thus connected) regions which we call linear regions. For any \tilde{y} in one of these regions, f is described by a unique local affine function h_j . The number of those regions is equal to the number of affine pieces of f .

V. FOLDING-BASED NEURAL DECODING OF A_n

In this section, we first prove that the lattice basis from the $n \times n$ Gram matrix in (4) is VR (all bases are equivalent modulo rotations and reflections). We count the number of pieces of the decision boundary function. We then build a deep ReLU network which computes efficiently this function via *folding*. Finally, we use the fact that the $n-2$ -dimensional hyperplanes partitioning \mathcal{D} are *not* in “general position” in order to prove that a two-layer network needs an exponential number of neurons to compute the function.

Consider a basis for the lattice A_n with all vectors from the first lattice shell. Also, the angle between any two basis vectors is $\pi/3$. Let J_n denote the $n \times n$ all-ones matrix and I_n the identity matrix. The Gram matrix is

$$\Gamma = GG^T = J_n + I_n. \quad (4)$$

Theorem 3. A lattice basis defined by the Gram matrix (4) is Voronoi-reduced.

See [15] - VII-C for the proof. Consequently, discriminating a point in $\mathcal{P}(\mathcal{B})$ with respect to the decision boundary leads to an optimal decoder.

A. Number of pieces of the decision boundary function

We count the number of pieces, and thus linear regions, of the decision boundary function f . We start with the following lemma involving i -simplices.

Lemma 1. *Consider an A_n -lattice basis defined by the Gram matrix (4). The decision boundary function f has a number of affine pieces equal to*

$$\sum_{i=0}^n i \times (\# \text{ regular } i\text{-simplices}), \quad (5)$$

where, for each i -simplex, only one corner x belongs to $\mathcal{C}_{\mathcal{P}(\mathcal{B})}^1$ and the other corners constitute the set $\tau_f(x) \cap \mathcal{C}_{\mathcal{P}(\mathcal{B})}^0$.

Proof. A key property of this basis is

$$\begin{aligned} \forall x \in \mathcal{C}_{\mathcal{P}(\mathcal{B})}^0, x' \in A_n \setminus \{b_j, 0\}, 2 \leq j \leq n : \\ x + b_j \in \tau_f(x + b_1), x + x' \notin \tau_f(x + b_1) \cap \mathcal{C}_{\mathcal{P}(\mathcal{B})}^0. \end{aligned} \quad (6)$$

It is obvious that $\forall x \in \mathcal{C}_{\mathcal{P}(\mathcal{B})}^0: x + b_1 \in \mathcal{C}_{\mathcal{P}(\mathcal{B})}^1$. This implies that any given point $x \in \mathcal{C}_{\mathcal{P}(\mathcal{B})}^0$ and its neighbors $\tau_f(x) \cap \mathcal{C}_{\mathcal{P}(\mathcal{B})}^0$ form a regular simplex \mathcal{S} of dimension $|\tau_f(x) \cap \mathcal{C}_{\mathcal{P}(\mathcal{B})}^0|$. This clearly appears on Figure 2. Now, consider the decision boundary function of a k -simplex separating the top corner (i.e. $\mathcal{C}_{\mathcal{S}}^1$) from all the other corners (i.e. $\mathcal{C}_{\mathcal{S}}^0$). This function is convex and has k pieces. The maximal dimensionality of such a simplex is obtained by taking the points $0, b_1$, and the $n - 1$ points $b_j, j \geq 2$. \square

Theorem 4. *Consider an A_n -lattice basis defined by the Gram matrix (4). The decision boundary function f has a number of affine pieces equal to*

$$\sum_{i=1}^n i \cdot \binom{n-1}{n-i}. \quad (7)$$

Proof. From Lemma 1, what remains to be done is to count the number of k -simplices. We walk in $\mathcal{C}_{\mathcal{P}(\mathcal{B})}^0$ and for each of the 2^{n-1} points $x \in \mathcal{C}_{\mathcal{P}(\mathcal{B})}^0$ we investigate the dimensionality of the simplex where the top corner is $x + b_1 \in \mathcal{C}_{\mathcal{P}(\mathcal{B})}^1$. This is achieved by counting the number of elements in $\tau_f(x + b_1) \cap \mathcal{C}_{\mathcal{P}(\mathcal{B})}^0$, via the property given by (6). Starting from the origin, one can form a n -simplex with $0, b_1$, and the $n - 1$ other basis vectors. Then, from any $b_{j_1}, j_1 \neq 1$, one can only add the $n - 1$ remaining basis vectors to generate a simplex in $\mathcal{P}(\mathcal{B})$. Indeed, if we add again b_{j_1} , the point goes outside $\mathcal{P}(\mathcal{B})$. Hence, we get a $n - 1$ -simplex and there are $\binom{n-1}{1}$ ways to choose b_{j_1} : any basis vectors except b_1 . Similarly, if one starts the simplex from $b_{j_1} + b_{j_2}$, one can form a $n - 2$ -simplex in $\mathcal{P}(\mathcal{B})$ and there are $\binom{n-1}{2}$ ways to choose $b_{j_1} + b_{j_2}$. In general, there are $\binom{n-1}{k}$ ways to form a $n - k$ -simplex. Applying the previous lemma and summing over $k = n - i = 0 \dots n - 1$ gives the announced result. \square

B. Decoding via folding

Obviously, at a location \tilde{y} , we do not want to compute all affine pieces in (3) whose number is given by (7) in order to evaluate f . To reduce the complexity of this evaluation, the idea is to exploit the symmetries of f by ‘‘folding’’ the function and mapping distinct regions of the input domain to the same location. If folding is applied sequentially, i.e. fold a region that has already been folded, it is easily seen that the gain becomes exponential. The notion of folding the input space in the context of neural networks was introduced in [9].

Given the basis orientation as in Theorem 1, the projection of b_j on \mathcal{D} is b_j itself, for $j \geq 2$. We also denote the bisector hyperplane between two vectors b_j, b_k by $BH(b_j, b_k)$ and its normal vector is taken to be $v_{j,k} = b_j - b_k$. We define the folding transformation $F: \mathcal{D} \rightarrow \mathcal{D}'$ as follows: let $\tilde{y} \in \mathcal{D}$, for all $2 \leq j < k \leq n$, compute $\tilde{y} \cdot v_{j,k}$ (the first coordinate of $v_{j,k}$ is zero). If the scalar product is non-positive, replace \tilde{y} by its mirror image with respect to $BH(b_j, b_k)$. There exist $(n - 1)(n - 2)/2$ hyperplanes for mirroring.

Theorem 5. *Let us consider the lattice A_n defined by the Gram matrix (4). We have (i) $\mathcal{D}' \subset \mathcal{D}$, (ii) for all $\tilde{y} \in \mathcal{D}$, $f(\tilde{y}) = f(F(\tilde{y}))$ and (iii) f has exactly*

$$2n + 1 \quad (8)$$

pieces on \mathcal{D}' . This is to be compared with (7).

See [15] - VII-D for the proof.

Example 1 (Continued). *The function f restricted to \mathcal{D}' (i.e. the function to evaluate after folding), say $f_{\mathcal{D}'}$, is*

$$f_{\mathcal{D}'} = \left[h_{p1} \vee h_1 \right] \wedge \left[h_{p2} \vee h_2 \right] \wedge \left[h_{p3} \right]. \quad (9)$$

The general expression of $f_{\mathcal{D}'}$ for any dimension is available in [15] - VII-A.

C. From folding to a deep ReLU network

For sake of simplicity and without loss of generality, in addition to the standard ReLU activation function $\text{ReLU}(a) = \max(0, a)$, we also allow the function $\max(0, -a)$ and the identity as activation functions in the network.

To implement a reflection, one can use the following strategy. Step 1: rotate the axes to have the j -th axis e_j perpendicular to the reflection hyperplane and shift the point (i.e. the j -th coordinate) to have the reflection hyperplane at the origin. Step 2: take the absolute value of the j -th coordinate. Step 3: do the inverse operation of step 1.

Now consider the ReLU network illustrated in Figure 3. The edges between the input layer and the hidden layer represent the rotation matrix, where the j -th column is repeated twice, and p is a bias applied on the j -th coordinate. Within the dashed square, the absolute value of the j -th coordinate is computed and shifted by $-p$. Finally, the edges between the hidden layer and the output layer represent the inverse rotation matrix. This ReLU network computes a reflection. We call it a reflection block.

All reflections can be naively implemented by a simple concatenation of reflection blocks. This leads to a very deep




Fig. 3. Reflection ReLU network (called reflection block).

and narrow network of depth $\mathcal{O}(n^2)$ and width $\mathcal{O}(n)$.

Regarding the $2n+1$ remaining pieces after folding, we have two options (in both cases, the number of operations involved is negligible compared to the previous folding operations). To directly discriminate the point with respect to f , we implement the HLD on these remaining pieces with two additional hidden layers (see e.g. Figure 2 in [5]): project y_{folded} on the $2n+1$ hyperplanes (with one layer of width $2n+1$) and compute the associated Boolean equation with an additional hidden layer. If needed, we can alternatively evaluate $f(\tilde{y})$ via $\mathcal{O}(\log(n))$ additional hidden layers. First, compute the $n-1$ $2-\vee$ via two layers of size $\mathcal{O}(n)$ containing several “max ReLU networks” (see e.g. Figure 3 in [2]). Then, compute the $n-\wedge$ via $\mathcal{O}(\log(n))$ layers. Note that $f(\tilde{y})$ can also be used for discrimination via the sign of $y^i - f(\tilde{y})$.

The total number of parameters in the whole network is $\mathcal{O}(n^4)$. In [15] - VII-G, we quickly discuss whether or not this can be improved.

Eventually, the CPV is solved by using n such networks in parallel (this could also be optimized). The final network has width $\mathcal{O}(n^2)$ and depth $\mathcal{O}(n^2)$.

D. Decoding via a shallow network

A two-layer ReLU network with n inputs and w_1 neurons in the hidden layer can compute a CPWL function with at most $\sum_{j=0}^n \binom{w_1}{j}$ pieces [10]. This result is easily understood by noticing that the non-differentiable part of $\max(0, a)$ is a $n-2$ -dimensional hyperplane that separates two linear regions. If one sums w_1 functions $\max(0, d_j \cdot y)$, where $d_j, 1 \leq j \leq w_1$, is a random vector, one gets w_1 of such $n-2$ -hyperplanes. The rest of the proof consists in counting the number of linear regions that can be generated by these w_1 hyperplanes. The number provided by the previous formula is attained if and only if the hyperplanes are in general position. Clearly, in our situation the $n-2$ -hyperplanes partitioning \mathcal{D} are not in general position: the hyperplane arrangement is not *simple*. The proof of the following theorem, available in [15] - VII-E, consists in finding a lower bound on the number of such $n-2$ -hyperplanes.

Theorem 6. *A ReLU network with one hidden layer needs at least*

$$\sum_{i=2}^n (i-1) \cdot \binom{n-1}{n-i} \quad (10)$$

neurons to solve the CVP for the lattice A_n .

VI. CONCLUSIONS

We recently applied this theory to a SVR basis of lattices E_n , $6 \leq n \leq 8$. The decision boundary function has a number of pieces equal to

$$\sum_{i=3}^{n-1} \binom{n-3}{n-i} \left(1 + (n-i) + 2 \left[1 + 2(n-i) + \binom{n-i}{2} \right] + \left[1 + 3(n-i) + 3 \binom{n-i}{2} + \binom{n-i}{3} \right] \right) + 1,$$

a number which we successfully linearized via folding.

From a learning perspective, our findings suggest that many optimal decoders may be contained in the function class Φ of deep ReLU networks. Learnability results of the restricted model [1][3] show that the sample complexity is then $m_{\Phi}(\epsilon, \delta) = \mathcal{O}(WL \log(W)/\epsilon)$, avoiding the $1/\epsilon^2$ of the general model (where W is the number of parameters in the network and L the number of layers).

Additionally, the folding approach suits very well the non-uniform finite sample bounds of the information bottleneck framework [13]. Indeed, if we model the input of the network and its i -th layer after the i -th reflection block by random variables Y and Y_i , clearly $I(Y; Y_i)$ is reduced compared to $I(Y; Y_{i-1})$ for any distribution of Y .

REFERENCES

- [1] M. Anthony, P. Bartlett. *Neural Network Learning: Theoretical Foundations*. Cambridge University Press, 2009.
- [2] R. Arora, A. Basu, P. Mianjy, and A. Mukherjee “Understanding deep neural networks with rectified linear units.” *International Conference on Learning Representations*, 2018.
- [3] P. Bartlett, N. Harvey, C. Liaw, and A. Mehrabian, (2017). “Nearly-tight VC-dimension and pseudodimension bounds for piecewise linear neural networks.” arXiv preprint arXiv:1703.02930, Oct. 2017.
- [4] J. Conway and N. Sloane. *Sphere packings, lattices and groups*. Springer-Verlag, New York, 3rd edition, 1999.
- [5] V. Corlay, J.J. Boutros, P. Ciblat, and L. Brunel, “Neural Lattice Decoders,” arXiv preprint arXiv:1807.00592, July. 2018.
- [6] H. Coxeter. *Regular Polytopes*. Dover, NY, 3rd ed., 1973.
- [7] R. Eldan and O. Shamir. “The power of depth for feedforward neural networks.” *29th Annual Conference on Learning Theory*, pp. 907–940, 2016.
- [8] J. Hästad. “Almost optimal lower bounds for small depth circuits.” *Proc. 18th ACM symposium on Theory of computing*, pp. 6–20, 1986.
- [9] G. Montufar, R. Pascanu, K. Cho, and Y. Bengio, “On the Number of Linear Regions of Deep Neural Networks,” *Advances in neural information processing systems*, pp. 2924–2932, 2014.
- [10] R. Pascanu, G. Montufar, and Y. Bengio. “On the number of inference regions of deep feed forward with piece-wise linear activations,” arXiv preprint arXiv:1312.6098, Dec. 2013.
- [11] M. Raghu, B. Poole, J. Kleinberg, S. Ganguli, and J. Sohl-Dickstein. “On the expressive power of deep neural networks,” arXiv preprint arXiv:1606.05336, Jun. 2016.
- [12] I. Safran and O. Shamir. “Depth-width tradeoffs in approximating natural functions with neural networks.” *In International Conference on Machine Learning*, pp. 2979–2987, 2017
- [13] O. Shamir, S. Sabato, and N. Tishby, “Learning and generalization with the information bottleneck,” *Theor. Comput. Sci.*, vol. 411, no.29-30, pp. 2696–2711, 2010.
- [14] M. Telgarsky. “Benefits of depth in neural networks,” *29th Annual Conference on Learning Theory*, pp. 1517–1539, 2016.
- [15] The long version of this paper including an appendix is available at: www.josephboutros.org/isit2019lattices.pdf.