

Revisiting the Expressive Power of Simple ReLU Neural Networks

Hirotada Honda

Faculty of Information Networking for Innovation and Design

Toyo University

Tokyo, Japan

honda.hirotada@iniad.org

Abstract—This paper investigates the expressive power of shallow ReLU neural networks from a geometric and combinatorial viewpoint. While the approximation capabilities of ReLU networks have been studied extensively, far less attention has been paid to the geometric structure of the decision boundaries that such networks induce. Focusing on networks with a single hidden layer acting on the 2-dimensional Euclidean space, we analyze how ReLU activations partition the input space into polyhedral regions and how these regions combine to form piecewise-linear decision boundaries. Using tools from the theory of affine hyperplane arrangements, we provide a simple characterization of the bending complexity of decision boundaries and relate it directly to the chamber structure of the underlying arrangement. This geometric perspective offers a transparent description of the mechanisms by which even small ReLU networks generate nonlinear separators, and serves as a foundation for future investigations into deeper or higher-dimensional architectures.

Index Terms—ReLU, hyperplane arrangement, expressive power, oriented matroid.

I. INTRODUCTION

Rectified Linear Unit (ReLU) neural networks have become a central model in modern machine learning, owing to their computational efficiency and strong empirical performance. From a theoretical perspective, the expressive power of ReLU networks has been studied extensively, with a particular focus on approximation properties, representation capacity, and depth-width trade-offs. Various works have shown that ReLU networks can approximate broad classes of functions and that network depth plays a crucial role in enhancing expressive efficiency.

In this paper, we revisit the expressive power of simple ReLU networks from a geometric viewpoint. Whereas much existing literature studies approximation accuracy or function classes, our interest lies instead in the *geometry of decision boundaries* induced by such networks. Because ReLU networks are piecewise-linear, their decision boundaries consist of linear segments joined at bending points. The position and number of these bends encode essential aspects of the network’s expressive power.

Research questions. Motivated by the geometric nature of ReLU networks, we pose the following questions:

- (i) *What kinds of geometric decision regions (e.g., shapes of label-1 sets) can be realized by ReLU networks of fixed width and depth?*

- (ii) *How many hidden units and layers are required to realize a target decision region, or to approximate it to arbitrary precision?*

These questions are fundamental for understanding the structural limitations and capabilities of neural networks. They also complement classical approximation-theoretic results by focusing on the geometry of classification boundaries rather than the approximation of scalar-valued functions.

Our analysis begins with the simplest nontrivial case: ReLU networks with a single hidden layer acting on \mathbb{R}^2 . Despite this seemingly restrictive setting, the network already exhibits rich geometric behavior. The decision boundary is not a single line but a piecewise-linear curve composed of multiple linear segments connected at bending points. Understanding how these bends arise and how many such bends can occur is a key step toward answering the research questions above.

A central observation underlying our approach is that each ReLU unit induces a hyperplane in the input space corresponding to its activation boundary (pre-activation value zero). Consequently, the activation patterns of a shallow ReLU network correspond to the chambers of an associated affine hyperplane arrangement. Within each chamber, the network computes an affine function, and the decision boundary is obtained by stitching together the affine decision surfaces across adjacent chambers.

This viewpoint connects the study of ReLU networks directly with the classical theory of hyperplane arrangements and oriented matroids. In particular, our results show that the number of bending points of a connected component of the decision boundary is controlled by the number of chambers intersected by that component. Thus, the geometric complexity of the decision boundary is governed by combinatorial properties of the induced hyperplane arrangement.

Although the setting is intentionally minimal, it already reveals the structural mechanism by which ReLU networks generate nonlinear decision boundaries. Moreover, this perspective provides a principled way to reason about which decision regions can or cannot be realized using networks of a given size. We view this work as a first step toward a general geometric framework for classifying the expressive power of ReLU networks.

The remainder of the paper is organized as follows. Section II reviews affine hyperplane arrangements and oriented matroids. In Section III, we review some related works. Section IV presents our main result on the bending complexity of decision boundaries. Section V illustrates these ideas with explicit examples. Section VI discusses extensions, limitations, and open problems.

II. TERMS AND NOTATIONS

In this section, we introduce the terminology and notations that will be used throughout the paper. Our exposition is intentionally brief and focuses only on concepts that are directly relevant to the geometric analysis of decision boundaries of ReLU neural networks. In particular, we review affine hyperplane arrangements and oriented matroids from a geometric viewpoint.

A. Affine Hyperplane Arrangements

An *affine hyperplane* in \mathbb{R}^d is a subset of the form

$$H = \{x \in \mathbb{R}^d \mid a^\top x + b = 0\}, \quad (\text{II.1})$$

where $a \in \mathbb{R}^d \setminus \{0\}$ and $b \in \mathbb{R}$. A finite collection of affine hyperplanes

$$\mathcal{A} = \{H_1, \dots, H_n\}, \quad (\text{II.2})$$

is called an *affine hyperplane arrangement*.

An arrangement \mathcal{A} decomposes \mathbb{R}^d into a finite number of connected open regions, referred to as *chambers*. Each chamber is a convex polyhedron, and the union of all chambers, together with their lower-dimensional faces, provides a polyhedral cell decomposition of the ambient space. In the planar case $d = 2$, which is the primary focus of this paper, chambers are polygonal regions bounded by line segments.

Affine hyperplane arrangements arise naturally in the study of ReLU neural networks. Each ReLU unit is associated with a pre-activation function of the form $a^\top x + b$, and the corresponding hyperplane $a^\top x + b = 0$ separates the input space into regions where the activation is either active or inactive. As a result, the collection of all pre-activation hyperplanes induces an arrangement whose chambers characterize the activation patterns of the network.

Within each chamber of the arrangement, the ReLU neural network reduces to an affine function, since the activation status of every ReLU unit is fixed. Consequently, the decision boundary of the network intersects each chamber along an affine subspace of codimension one, typically a line segment in \mathbb{R}^2 .

For a comprehensive treatment of affine hyperplane arrangements and their combinatorial and geometric properties, we refer to the classical monograph by Orlik and Terao [1].

B. Oriented Matroids

Oriented matroids provide a combinatorial framework for encoding the structure of affine hyperplane arrangements. Given an arrangement

$$\mathcal{A} = \{H_1, \dots, H_n\}, \quad H_i = \{x \mid a_i^\top x + b_i = 0\}, \quad (\text{II.3})$$

one associates to each point $x \in \mathbb{R}^d$ a sign vector

$$\begin{aligned} \text{sign}_{\mathcal{A}}(x) &= (\text{sign}(a_1^\top x + b_1), \dots, \text{sign}(a_n^\top x + b_n)) \\ &\in \{+, 0, -\}^n. \end{aligned} \quad (\text{II.4})$$

The collection of all such sign vectors forms the set of *covectors* of the oriented matroid associated with the arrangement. Covectors with no zero entries correspond to chambers, while covectors with at least one zero entry correspond to lower-dimensional faces lying on one or more hyperplanes.

A systematic and authoritative exposition of oriented matroids and their deep connections to affine hyperplane arrangements is given in the monograph by Björner et al. [2].

From this perspective, the combinatorial data of the oriented matroid captures how hyperplanes intersect and how the arrangement partitions the ambient space, independently of the specific geometric realization. This abstraction is particularly useful when analyzing piecewise-linear structures, such as those arising from ReLU activations, where the geometry is determined by sign patterns of affine functions.

In the context of ReLU neural networks, the activation pattern of the hidden units at an input point x naturally defines a covector of the associated oriented matroid. The decision boundary, given by the zero set of the network output, corresponds to a collection of faces where additional linear constraints are satisfied. This viewpoint allows us to relate geometric properties of decision boundaries, such as bending points, to purely combinatorial features of hyperplane arrangements and their oriented matroids.

The relevance of hyperplane arrangements to neural networks has also been highlighted in studies on linear region complexity and expressivity of ReLU networks; see, e.g., Raghu et al. [3].

III. RELATED WORK

The expressive power of neural networks has been studied from various perspectives for several decades.

A classical result by Leshno et al. [4] established the universal approximation property of feedforward neural networks with non-polynomial activation functions, including the ReLU activation. Earlier foundational contributions include the seminal works of Cybenko [5] and Funahashi [6], [7], as well as studies on the capability of shallow perceptron architectures [8]. Since then, numerous works have investigated the approximation capabilities of ReLU neural networks, with particular emphasis on depth-width trade-offs and representation efficiency. For a general early overview of neural network models and learning mechanisms, see the introductory article by Lippmann [9].

More recently, several studies have focused on the geometric and combinatorial properties of ReLU networks. It is well known that ReLU activations induce piecewise-linear functions, and that the input space is partitioned into polyhedral regions on which the network behaves affinely. This observation has motivated analyses based on region counting and linear region complexity, as well as connections to hyperplane arrangements and polyhedral geometry.

From a modern approximation-theoretic viewpoint, neural networks can be understood as nonlinear approximation schemes whose expressive power is closely tied to compositional structure and sparsity; see, for example, the survey by DeVore, Hanin, and Petrova [10]. Related work has analyzed the role of compositional depth in enhancing nonlinear approximation efficiency [11], as well as connections between deep residual networks and nonlinear control systems [12].

A fundamental line of theoretical work on ReLU neural networks concerns the combinatorial complexity induced by their piecewise-linear activations. Because each ReLU unit introduces a linear threshold, a feedforward ReLU network partitions the input space into finitely many polyhedral regions, commonly referred to as *response regions*. On each such region, the network realizes an affine map.

Pascanu et al. [13] provided one of the earliest systematic analyses of this phenomenon, deriving upper and lower bounds on the number of response regions as a function of network depth and width. Their results demonstrate that network depth plays a crucial role in exponentially increasing the number of linear regions, thereby offering a quantitative explanation for the expressive efficiency of deep architectures.

Subsequent works have refined these bounds and further clarified the relationship between network architecture and region counts, often through connections to hyperplane arrangements and combinatorial geometry. While this line of research primarily focuses on *counting* response regions, it leaves open more detailed questions concerning the *geometric and topological structure* of individual regions and their unions, such as convexity, connectivity, and the global organization of decision boundaries.

These works interpret ReLU networks as generating semialgebraic sets and investigate structural properties arising from such representations. Although our focus is different, both approaches highlight the fundamental role played by piecewise-linear and combinatorial structures induced by ReLU activations.

However, despite these advances, comparatively few works have examined how the geometry of *decision boundaries* themselves can be characterized in a precise and explicit manner. In particular, the extent to which the decision boundary of a target population can be represented by a ReLU neural network, and how its geometric complexity depends on the induced polyhedral structure, remains less understood.

The present work contributes to this line of research by providing a geometric and combinatorial description of decision boundaries for shallow ReLU neural networks in low dimensions. By explicitly relating bending points of decision boundaries to chambers of an associated affine hyperplane arrangement, we offer a transparent framework for understanding the piecewise-linear structure induced by ReLU activations.

IV. MAIN RESULT

In this section, we present our main result on the bending complexity of decision boundaries arising from shallow ReLU

networks in two dimensions.

Theorem IV.1 (Bending complexity of shallow ReLU decision boundaries). *Let $f : \mathbb{R}^2 \rightarrow \mathbb{R}$ be the decision function of a ReLU neural network with a single hidden layer of width two and a linear readout. Write the pre-activation as $z = Wx + b \in \mathbb{R}^2$, and assume $\text{rank}(W) = 2$ so that the induced hyperplanes*

$$H_1 := \{z_1 = 0\}, \quad H_2 := \{z_2 = 0\}, \quad (\text{IV.1})$$

form a rank-2 affine hyperplane arrangement in \mathbb{R}^2 . Then the decision boundary

$$\{x \in \mathbb{R}^2 \mid f(x) = 0\}, \quad (\text{IV.2})$$

is a piecewise-linear curve such that each connected component is contained in at most three chambers of the arrangement induced by $\{H_1, H_2\}$.

Consequently, any connected component of the decision boundary has at most two bending points.

A. Proof of Theorem IV.1

Since $\text{rank}(W) = 2$, the affine map

$$z = Wx + b : \mathbb{R}^2 \rightarrow \mathbb{R}^2, \quad (\text{IV.3})$$

is a bijection. The decision boundary in the x -space,

$$\{x \in \mathbb{R}^2 \mid f(x) = 0\}, \quad (\text{IV.4})$$

is therefore the affine preimage of the curve Γ in the z -space. An affine bijection preserves piecewise-linearity as well as the number of bending points of each connected component. Hence it suffices to analyze the geometry of Γ in the z -plane.

By expanding the definition of the ReLU, we can write the decision function in z -coordinates as

$$y(z) = w_1\sigma(z_1) + w_2\sigma(z_2) + c \\ = \begin{cases} w_1z_1 + w_2z_2 + c, & z_1 > 0, z_2 > 0, \\ w_2z_2 + c, & z_1 \leq 0, z_2 > 0, \\ w_1z_1 + c, & z_1 > 0, z_2 \leq 0, \\ c, & z_1 \leq 0, z_2 \leq 0. \end{cases} \quad (\text{IV.5})$$

Thus, on each open quadrant the restriction of y is an affine function of (z_1, z_2) . Consequently, the intersection of the decision boundary $\Gamma = \{y(z) = 0\}$ with each open quadrant is either empty or a single line segment. Any bending point of Γ can occur only when Γ crosses one of the coordinate axes $z_1 = 0$ or $z_2 = 0$, where the affine expression of y changes.

We next show that Γ intersects each coordinate axis at most once.

Claim. The curve Γ intersects the axes $\{z_1 = 0\}$ and $\{z_2 = 0\}$ at most once each.

Proof of the claim. Suppose that $c < 0$ holds. We treat the z_2 -axis; the argument for the z_1 -axis is analogous.

On the z_2 -axis, we have $z_1 = 0$, hence

$$y(0, z_2) = w_1\sigma(0) + w_2\sigma(z_2) + c = w_2\sigma(z_2) + c. \quad (\text{IV.6})$$

For $z_2 > 0$ (the upper half-axis), this is the affine function $w_2z_2 + c$, which has at most one zero. For $z_2 < 0$ (the lower

half-axis), we have $\sigma(z_2) = 0$, so $y(0, z_2) = c < 0$ and thus there is no zero at all. Therefore Γ can intersect the z_2 -axis in at most one point. We can consider the case of $c > 0$ in a similar manner. The same reasoning with z_1 and z_2 interchanged shows that Γ intersects the z_1 -axis in at most one point. \square

Combining these observations, we obtain the desired bound on the number of bending points. Each connected component of Γ is contained in a union of quadrants and may cross from one quadrant to another only by crossing one of the axes. Thus, Γ never enters at least one quadrant and intersects each axis at most once, and a connected component can visit at most three quadrants. Moreover, every time the component passes from one quadrant to another, the local defining affine function of y changes, and the decision boundary acquires a bending point.

Hence, a connected component of Γ can have at most two such transitions between quadrants, and therefore at most two bending points. Since the decision boundary in the x -space is the affine preimage of Γ , the same bound holds for its connected components. This completes the proof of Theorem IV.1. \square

V. SIMPLE EXPERIMENTS

In this section, we illustrate our theoretical results using a simple numerical example. We consider a ReLU neural network with a single hidden layer of width two and a linear readout. The network is defined on \mathbb{R}^2 by

$$y(x) = \sigma(z_1(x)) + \sigma(z_2(x)) - 3, \quad (\text{V.1})$$

where

$$z_1(x) = 2x_1 + x_2 + 3, \quad z_2(x) = x_1 + 2x_2 + 1. \quad (\text{V.2})$$

The affine hyperplanes $z_1 = 0$ and $z_2 = 0$ partition the input space into four chambers, corresponding to the sign patterns

$$(+, +), (+, -), (-, +), (-, -).$$

Within each chamber, the ReLU activations are fixed, and the network output reduces to an affine function. The correspondence between chambers, local expressions of y , and the associated boundary pieces is summarized in Table I.

TABLE I

LOCAL AFFINE REPRESENTATIONS OF THE NETWORK OUTPUT ON EACH CHAMBER. THE DECISION BOUNDARY $y = 0$ EXISTS ONLY IN CHAMBERS WHERE A NONCONSTANT AFFINE EXPRESSION IS OBTAINED.

Chamber ($\text{sign}(z_1), \text{sign}(z_2)$)	Local form of y	Boundary piece
$(+, +)$	$z_1 + z_2 - 3$	Line segment
$(+, -)$	$z_1 - 3$	Line segment
$(-, +)$	$z_2 - 3$	Line segment
$(-, -)$	Constant -3	None

As a consequence, the decision boundary intersects exactly three chambers and is composed of three connected linear segments. Therefore, the boundary has two bending points, in accordance with the theory developed in the previous sections.

The resulting decision boundary is visualized in Fig. 1, where the piecewise linear structure and the two bending points can be clearly observed.




Fig. 1. Figure X: Decision boundary of a ReLU neural network with a single hidden layer of width two. The boundary consists of three linear segments connected by two bending points, occurring at transitions between chambers of the induced hyperplane arrangement.

Similarly to the width-two case, the width-three example in Figure 2 illustrates that each linear segment of the decision boundary corresponds to a fixed sign pattern of the pre-activations (z_1, z_2, z_3) , while each bending point arises when the input crosses one of the hyperplanes $\{z_i = 0\}$ that bound the chambers of the induced hyperplane arrangement.

VI. DISCUSSION AND CONCLUSION

The analysis presented in this paper focuses on shallow ReLU neural networks in low dimensions, where the geometry of decision boundaries can be described explicitly. In more general settings, however, a deeper understanding requires a more systematic use of tools from the theory of hyperplane arrangements.

In particular, the combinatorial complexity of decision boundaries is closely related to the chamber structure induced by the pre-activation hyperplanes. For higher width or higher-dimensional input spaces, explicitly counting the number of chambers becomes nontrivial. In such cases, classical results from hyperplane arrangement theory, including those based on characteristic polynomials, may provide useful bounds on the number of possible chambers and, consequently, on the maximal complexity of decision boundaries.

Another important direction is the extension to deeper networks. Each additional hidden layer effectively refines the partition of the input space, leading to a hierarchical structure of polyhedral decompositions. Understanding how these refinements interact across layers, and how they affect

Example decision boundary with 3-node hidden ReLU layer




Fig. 2. Decision boundary of a ReLU neural network with a single hidden layer of width three. The decision boundary is a piecewise-linear curve composed of three linear segments joined at four bending points, which occur exactly at the transitions between different chambers of the underlying hyperplane arrangement in the hidden pre-activation space.

the geometry and topology of decision boundaries, remains an open and challenging problem.

Finally, although our experiments are restricted to low-dimensional examples, the underlying combinatorial perspective suggests that qualitative features of ReLU decision boundaries—such as bending points, connected components, and overall piecewise-linear complexity—are governed by structural properties of the induced hyperplane arrangements rather than by specific numerical parameters. We believe that further exploration along these lines may provide a fruitful bridge between classical combinatorial geometry and modern neural network theory.

VII. USE OF LARGE LANGUAGE MODELS (LLMs)

We used ChatGPT for language polishing of author-written text. All results were reviewed, run, and validated by the author. All scientific ideas, model designs, proofs, and claims are the authors' work; the LLM is not an author.

ACKNOWLEDGMENT

This work was supported by Inoue Enryō Memorial Grant, Toyo University.

REFERENCES

- [1] P. Orlik and H. Terao, *Arrangements of Hyperplanes*. Springer, 1992.
- [2] A. Björner, M. Las Vergnas, B. Sturmfels, N. White, and G. M. Ziegler, *Oriented Matroids*, 2nd ed. Cambridge University Press, 1999.
- [3] M. Raghu, B. Poole, J. Kleinberg, S. Ganguli, and J. Dickstein, On the expressive power of deep neural networks, in *Proc. International Conference on Machine Learning (ICML)*, 2017.
- [4] M. Leshno, V. Y. Lin, A. Pinkus, and S. Schocken, Multilayer feedforward networks with a nonpolynomial activation function can approximate any function, *Neural Networks*, vol. 6, no. 6, pp. 303–314, 1993. DOI: 10.1016/S0893-6080(05)80131-5.
- [5] G. Cybenko, Approximation by superpositions of a sigmoidal function, *Mathematics of Control, Signals and Systems*, vol. 2, no. 4, pp. 303–314, 1989. DOI: 10.1007/BF02551274.
- [6] K. Funahashi, On the approximate realization of continuous mappings by neural networks, *Neural Networks*, vol. 2, no. 3, pp. 183–192, 1989. DOI: 10.1016/0893-6080(89)90003-8.
- [7] K. Funahashi and Y. Nakamura, Neural networks, approximation theory, and dynamical systems (structure and bifurcation of dynamical systems), *Suuri-Kaiseki Kenkyujo Kokyuroku*, pp. 18–37, 1992. [Online]. Available: <http://hdl.handle.net/2433/82914>.
- [8] B. Irie and S. Miyake, Capabilities of three-layered perceptrons, in *Proc. IEEE International Conference on Neural Networks*, pp. 641–648, 1988. DOI: 10.1109/ICNN.1988.23901.
- [9] R. P. Lippmann, An introduction to computing with neural nets, *IEEE ASSP Magazine*, vol. 4, no. 2, pp. 4–22, 1987. DOI: 10.1109/MASSP.1987.1165576.
- [10] R. DeVore, B. Hanin, and G. Petrova, Neural network approximation, *Acta Numerica*, vol. 30, pp. 327–444, 2021. DOI: 10.1017/S0962492921000052.
- [11] Z. Shen, H. Yang, and S. Zhang, Nonlinear approximation via compositions, *arXiv preprint arXiv:1902.10170*, 2019. [Online]. Available: <http://arxiv.org/abs/1902.10170>.
- [12] P. Tabuada and B. Gharesifard, Universal approximation power of deep residual neural networks via nonlinear control theory, in *Proc. International Conference on Learning Representations (ICLR)*, 2021. [Online]. Available: <https://openreview.net/forum?id=-IXhmY16R3M>.
- [13] Pascanu, R., et al., On the Number of Response Regions of Deep Feedforward Networks with Piecewise Linear Activations, *arXiv preprint arXiv:1312.6098*, 2014.
- [14] B. I. Yun, A neural network approximation based on a parametric sigmoidal function, *Mathematics*, vol. 7, no. 3, 2019. DOI: 10.3390/math7030262. [Online]. Available: <https://www.mdpi.com/2227-7390/7/3/262>.