

ON TRANSVERSALITY OF BENT HYPERPLANE ARRANGEMENTS AND THE TOPOLOGICAL EXPRESSIVENESS OF RELU NEURAL NETWORKS

J. ELISENDA GRIGSBY AND KATHRYN LINDSEY

ABSTRACT. Let $F : \mathbb{R}^n \rightarrow \mathbb{R}$ be a feedforward ReLU neural network. It is well-known that for any choice of parameters, F is continuous and piecewise affine-linear. We lay some foundations for a systematic investigation of how the *architecture* of F impacts the geometry and topology of its possible decision regions, $F^{-1}((-\infty, t))$ and $F^{-1}((t, \infty))$, for binary classification tasks. Following the classical progression for smooth functions in differential topology, we first define the notion of a *generic, transversal* ReLU neural network and show that almost all ReLU networks are generic and transversal. We then define a partially-oriented linear 1-complex in the domain of F and identify properties of this complex that yield an obstruction to the existence of bounded connected components of a decision region. We use this obstruction to prove that a decision region of a generic, transversal ReLU network $F : \mathbb{R}^n \rightarrow \mathbb{R}$ with a single hidden layer of dimension $n + 1$ can have no more than one bounded connected component.

1. INTRODUCTION

Neural networks have rapidly become one of the most widely-used tools in the machine learning toolkit. Unfortunately, despite—or, perhaps, because of—their spectacular success in applications, significant foundational questions remain. Of these, we believe many would benefit greatly from the direct attention of theoretical mathematicians, particularly those in the geometric topology, nonlinear algebra, and dynamics communities. An important goal of this paper and its sequels is to advertise some of these problems to those communities.

Recall that one can view a (trained) feedforward neural network as a particular type of function, $F : \mathbb{R}^n \rightarrow \mathbb{R}^m$, between Euclidean spaces. The inputs to the function are data feature vectors and the outputs are often used to answer m -class classification problems by partitioning the input space into *decision regions* according to which component of the function output is maximized at that point.

The main purpose of the present work is to present a framework for studying the question: *How does the architecture of a feedforward neural network constrain the topology of its decision regions?* Here, the *architecture* of a feedforward neural network refers simply to the dimensions of the hidden layers. The neural networks we consider here will be fully-connected ReLU networks without skip connections. The *topological expressiveness* of an

JEG was partially supported by Simons Collaboration grant 635578.

KL was partially supported by NSF grant number DMS-1901247.

architecture is the collection of possible homeomorphism types of decision regions that can appear as the parameters vary (cf. [4, 12]).

First: why should the machine learning community care about *topological expressiveness*?

Recall that a cornerstone theoretical result in the study of neural networks—for a variety of activation functions, including the widely-used ReLU function that is our focus here—is the Universal Approximation Theorem ([6, 16, 1]), which says that a sufficiently high-dimensional neural network can approximate any continuous function on a compact set to arbitrary accuracy. This is the version of *representational power* or *expressiveness* frequently cited by practitioners as a guarantee that feedforward neural networks can yield a solution to any data question one might throw at them.

Yet continuous functions can be quite poorly behaved, and certain classes of poorly behaved continuous functions are undesirable targets for learning. For example, functions with high Lipschitz constants and ones whose partial derivatives are highly variable with respect to the input direction lead to the easy production of adversarial examples and hence to potentially poor generalization to unseen data (cf. [5, 24, 18]).

It is important to remark at this point that homeomorphism is a very coarse equivalence relation. Two different decision regions can be homeomorphic and still have quite different geometric properties (shape, volume, etc.). However, homeomorphism is a good equivalence relation to consider on a first pass, because the most fundamental global features of the data are preserved by homeomorphism (number of connected components and higher homology groups, e.g.). The flip side of this observation is that if a particular architecture lacks the topological expressiveness to capture obvious features inherent in a well-sampled labeled data set, it is unlikely to generalize well to unseen data.

In a more practical direction, we believe topological complexity is the right lens through which to study the implicit bias of stochastic gradient descent, i.e., the behavior of neural network functions during training. One mystery is the unreasonable success of overparameterized networks—those for which the number of trainable parameters far exceeds the size of the data set—in learning functions with low training loss and good generalization to unseen data [3]. The classical bias-variance trade-off curve suggests that increasing the number of parameters should lead the model to learn high-complexity functions that overfit the training set; overparameterized networks defy this intuition. An examination of the relationship between the topological features of an architecture’s loss landscape [20] and of its realizable functions [21] may shed light on this phenomenon. We lay groundwork for this in [8].

Last, we remark that ReLU and other piecewise-linear (PL) activation functions are now the default choice in the community for regression problems, which in turn underlie almost every other machine learning problem. Unfortunately, the resulting PL neural network functions are much trickier to study than smooth ones, because the classical foundational work in differential topology takes considerable effort to translate to the PL setting. ReLU, in particular, often gives rise to quite degenerate functions. In the sequel, we argue that this is a feature and not a bug [7].

We focus here on the simplest case of feedforward neural networks, $F : \mathbb{R}^n \rightarrow \mathbb{R}$, with 1-dimensional output. Such a network is often used to answer *binary* (aka Yes/No) classification problems by choosing a threshold, $t \in \mathbb{R}$, and declaring the sublevel set of t to be the “N” decision region, the superlevel set of t to be the “Y” decision region, and the level set of t to be the decision boundary. That is:

$$\begin{aligned} N_F(t) &:= F^{-1}((-\infty, t)) \\ B_F(t) &:= F^{-1}(\{t\}) \\ Y_F(t) &:= F^{-1}((t, \infty)). \end{aligned}$$

Classical results in differential topology now tell us that if F were smooth, we could perturb F slightly to be Morse, and the indices and values of its critical points would then provide strong information about the topology of its decision regions. Although a ReLU neural network map F is not smooth, F will *typically*¹ (i.e., for almost all choices of parameters) be sufficiently well-behaved that the information needed to understand much of the topology of its decision regions is extractable directly from the neural network parameters.

We begin by reviewing some standard results in the theory of affine hyperplane arrangements and convex polyhedra, relying heavily on Grunert’s work [10] on polyhedral complexes and Hanin-Rolnick’s work [14] (see also [19, 22, 13]) generalizing the classical notion of *hyperplane arrangements* to so-called *bent hyperplane arrangements* (Definition 6.1). Note that the appearance of polyhedral complexes in the study of ReLU networks is well-known, and made explicit, e.g., in the relationship between ReLU neural networks with rational parameters and tropical rational functions [28]. Our contribution here is to formalize the relationship and open a path for applying classical ideas in differential topology to study decision regions. We summarize our main results below. More precise versions and proofs of these theorems appear in later sections as indicated.

Theorem 1. *Let $F : \mathbb{R}^n \rightarrow \mathbb{R}$ be a ReLU neural network map. F is continuous and affine-linear on the cells of a canonical realization of the domain, \mathbb{R}^n , as a polyhedral complex, $\mathcal{C}(F)$.*

Moreover, when F is *generic* (Definition 2.4) and *transversal* (Definition 8.1), we explicitly identify cells of this polyhedral complex with natural objects defined by Hanin-Rolnick:

Theorem 2. *Let $F : \mathbb{R}^n \rightarrow \mathbb{R}$ be a generic, transversal ReLU neural network map. The n -cells of the canonical polyhedral complex, $\mathcal{C}(F)$, are the closures of the activation regions (Definition 6.3) of F , and the $(n-1)$ -skeleton of $\mathcal{C}(F)$ is the bent hyperplane arrangement (Definition 6.1) of F .*

We slightly extend classical transversality results (Theorems 7 and 8) to obtain:

Theorem 3. *Almost all ReLU neural networks are generic and transversal.*

¹We’d like to use the word *generically* here, but the term *generic* is unavoidably used in a different context later in the paper (Definition 2.4).

Letting D denote the dimension of the parameter space of a neural network, we show that a parametrized family of neural networks (Definition 2.3) $\mathcal{F} : \mathbb{R}^{n_0} \times \mathbb{R}^D \rightarrow \mathbb{R}$ is *piecewise smooth* in the following sense:

Theorem 4. *Every parametrized family of ReLU neural networks \mathcal{F} is smooth on the complement of a codimension 1 algebraic set.*

Many of the key observations in Theorems 1–4 were proved in [14]. The theorems above place those results in a broader context. Once we've established these foundational results, we turn our attention to addressing some first questions about architecture's impact on topological expressiveness. We begin by using the framework developed above to recast and reprove the result of Beise-Cruz-Schröder [2], proved independently by Johnson [17] and Hanin-Sellke [15], that inspired this study:²

Theorem 5. *For any integer $n \geq 2$, let $F : \mathbb{R}^n \rightarrow \mathbb{R}$ be a ReLU neural network, all of whose hidden layers have dimension $\leq n$. Then for any decision threshold $t \in \mathbb{R}$, each of $Y_F(t)$, $B_F(t)$, and $N_F(t)$ is either empty or unbounded.*

We also have the following new application. See [20] for related results:

Theorem 6. *Let $F : \mathbb{R}^n \rightarrow \mathbb{R}^{n+1} \rightarrow \mathbb{R}$ be a ReLU neural network with input dimension n and a single hidden layer of dimension $n+1$. Each decision region of F associated to a transversal threshold can have no more than one bounded connected component.*

A crucial player in the proof of Theorem 6 is the 1–skeleton, $\mathcal{C}(F)_1$, of the polyhedral complex, $\mathcal{C}(F)$, which is naturally endowed with a partial orientation pointing in the direction in which F increases (Definition 9.12).

Note that the partial orientation data can be extracted directly from the weight matrices of the neural network using the chain rule (Lemma 9.11).

This paper is heavy on definitions and notation, since we pulled from a variety of sources to lay necessary foundations for a consistent and general theory. Some sections may therefore be safely skimmed on a first reading and referenced only as needed to understand the proofs of the main results. Sections 2 and 9 largely fall into this category. Similarly, in Section 7 we establish important results about parameterized neural network maps and prove Theorem 4, but nothing in this section is referenced elsewhere in the paper.

Sections 3 and 4 establish notation and key terminology. We do the bulk of the technical work in Section 5, where we establish necessary transversality results; Section 6, where we prove Theorem 1; and Section 8, where we prove Theorems 2 and 3. The new applications can be found in Section 10, where we reprove Theorem 5 and prove Theorem 6.

Acknowledgements. The authors would like to thank Jesse Johnson for proposing many of the questions at the heart of this investigation; Boris

²The arguments of Beise-Cruz-Schröder and Johnson hold for a large class of activation functions, including the ReLU activation function studied here. Hanin-Sellke prove the statement only for ReLU, but their paper has a broader scope.

Hanin for illuminating conversations during the *Foundations of Deep Learning* program at the Simons Institute for the Theory of Computing in the summer of 2019 and numerous follow-on discussions that also included David Rolnick and Atlas Wang; Jenna Rajchgot for very helpful discussions about algebraic geometry; Yaim Cooper and Jordan Ellenberg for inspiration and encouragement. We would also like to thank the anonymous reviewers for pointing out gaps (now fixed) in some arguments and for a number of excellent suggestions that greatly improved the exposition.

2. LAYER MAPS AND HYPERPLANE ARRANGEMENTS

In what follows, let

- $\text{ReLU} : \mathbb{R} \rightarrow \mathbb{R}$ denote the function $\text{ReLU}(x) := \max\{0, x\}$, and
- $\sigma : \mathbb{R}^n \rightarrow \mathbb{R}^n$ denote the function that applies ReLU to each coordinate.

Definition 2.1. Let $n_0 \in \mathbb{N}$. A neural network defined on \mathbb{R}^{n_0} with ReLU activation function on all hidden layers and one-dimensional output is a finite sequence of natural numbers n_1, \dots, n_m together with affine maps $A_i : \mathbb{R}^{n_i} \rightarrow \mathbb{R}^{n_{i+1}}$ for $i = 0, \dots, m$. This determines a function $F : \mathbb{R}^{n_0} \rightarrow \mathbb{R}$, which we call the associated neural network map, given by the composition

$$\mathbb{R}^{n_0} \xrightarrow{F_1=\sigma \circ A_1} \mathbb{R}^{n_1} \xrightarrow{F_2=\sigma \circ A_2} \mathbb{R}^{n_2} \xrightarrow{F_3=\sigma \circ A_3} \dots \xrightarrow{F_m=\sigma \circ A_m} \mathbb{R}^{n_m} \xrightarrow{G=A_{m+1}} \mathbb{R}^1.$$

Such a neural network is said to be of architecture $(n_0, \dots, n_m; 1)$, depth $m + 1$, and width $\max\{n_1, \dots, n_m, 1\}$. The k^{th} layer map of such a neural network is the composition $\sigma \circ A_k$ for $k = 1, \dots, m$ and is the map $G = A_k$ for $k = m + 1$.

Remark 2.2. Note that in Definition 2.1 the activation function on the final layer map is the Identity function, not σ . Accordingly, we use the notation G on the output layer map to distinguish it from the hidden layer maps, F_k .

An affine map $A : \mathbb{R}^n \rightarrow \mathbb{R}^m$ is specified by a weight matrix $W \in M_{m \times n}(\mathbb{R})$ and a bias vector $\vec{b} \in \mathbb{R}^m$, as follows. Let $(W|\vec{b})$ denote the $m \times (n + 1)$ matrix whose final column is \vec{b} . For each $\vec{x} \in \mathbb{R}^n$, let $\vec{x}' := (\vec{x}, 1)$ be the element of \mathbb{R}^{n+1} whose first n coordinates are those of x and whose last coordinate is 1. Thus $A(\vec{x}) = (W|\vec{b})\vec{x}'$.

For a network architecture $(n_0, \dots, n_m; 1)$, we will denote by

$$D := \sum_{i=0}^m (n_i + 1)n_{i+1}$$

the total dimension of the parameter space of neural networks of that architecture. When unclear from context, we will specify the architecture in the notation for the dimension: $D(n_0, \dots, n_m; 1)$.

Definition 2.3. Let $(n_0, \dots, n_m; 1)$ be a network architecture. The parameterized family of ReLU neural networks with architecture $(n_0, \dots, n_m; 1)$ is the map $\mathcal{F} : \mathbb{R}^{n_0} \times \mathbb{R}^D \rightarrow \mathbb{R}$ defined as follows. For each $\mathbf{s} \in \mathbb{R}^D$, $\mathcal{F}_{\mathbf{s}} : \mathbb{R}^{n_0} \times \{\mathbf{s}\} \rightarrow \mathbb{R}$ is the ReLU neural network map associated to the weights and biases given by \mathbf{s} .

Observe that for each row of W , the corresponding row $(W_i \mid b_i) \in \mathbb{R}^{n+1}$, of the augmented matrix $(W \mid b) \in M_{m \times (n+1)}(\mathbb{R})$ determines a set:

$$(1) \quad S_i := \{\vec{x} \in \mathbb{R}^n \mid (W_i \mid b_i) \cdot (\vec{x} \mid 1) = 0\} \subseteq \mathbb{R}^n.$$

An *ordered affine solution set arrangement* in \mathbb{R}^n is a finite ordered set, $\mathcal{S} = \{S_1, \dots, S_m\}$, where each S_i is the solution set to an affine-linear equation as described above in equation (1). If $W_i = \mathbf{0}$, we say S_i is *degenerate*. In this case S_i is empty if $b_i \neq 0$ and S_i is all of \mathbb{R}^n if $b_i = 0$. An ordered affine solution set arrangement, $\mathcal{S} = \{S_1, \dots, S_m\}$, is said to be *degenerate* if at least one S_i is degenerate and *nondegenerate* otherwise. \mathcal{S} is said to be *in general position* (aka *generic*) if, for all subsets $\{S_{i_1}, \dots, S_{i_p}\} \subseteq \mathcal{S}$, it is the case that $S_{i_1} \cap \dots \cap S_{i_p}$ is an affine-linear subspace of \mathbb{R}^n of dimension $n - p$, where a negative-dimensional intersection is understood to be empty. Note that generic implies nondegenerate.

If S_i is nondegenerate, it is a hyperplane, and we will denote it by H_i . In this case, $\mathbb{R}^n \setminus H_i$ has two connected components,

$$\begin{aligned} H_i^+ &:= \{\vec{x} \in \mathbb{R}^n \mid (W_i \mid b_i) \cdot (\vec{x} \mid 1) > 0\} \\ H_i^- &:= \{\vec{x} \in \mathbb{R}^n \mid (W_i \mid b_i) \cdot (\vec{x} \mid 1) < 0\}, \end{aligned}$$

which endows H_i with a co-orientation, pointing toward H_i^+ .

The ordered affine solution set arrangement associated to a layer map, $\mathbb{R}^n \rightarrow \mathbb{R}^m$, of a neural network is the set $\mathcal{S} = \{S_1, \dots, S_m\}$ as in equation 1. A layer map of a neural network is said to be *degenerate* if its associated affine solution set arrangement \mathcal{S} is degenerate, and *nondegenerate* otherwise. A layer map of a neural network is said to be *generic* if the corresponding affine solution set arrangement is generic, and *nongeneric* otherwise.

If a solution set arrangement is nondegenerate, then each hyperplane, $H_i = S_i$, is co-oriented as described above, and we can therefore associate an *ordered, co-oriented hyperplane arrangement*, $\mathbf{A} = \{\mathbf{H}_1, \dots, \mathbf{H}_m\}$, in \mathbb{R}^n to the layer. By forgetting the ordering of the set and the co-orientations of the affine hyperplanes we obtain a classical hyperplane arrangement; that is, a finite set, $\mathcal{A} = \{H_1, \dots, H_m\}$, of affine hyperplanes in \mathbb{R}^n . We shall use \mathbf{A} (resp., \mathbf{H}_i) if the hyperplanes are equipped with a co-orientation and \mathcal{A} (resp., H_i) if not.

Definition 2.4. *A neural network whose layer maps are all nondegenerate is said to be nondegenerate. A neural network whose layer maps are all generic is said to be generic.*

The *rank* of a hyperplane arrangement \mathcal{A} in \mathbb{R}^n is the dimension of the space spanned by the normals to the hyperplanes in \mathcal{A} . (For more on the geometry and combinatorics of hyperplane arrangements, see [23].)

The following lemmas (cf. [23]) follow in a straightforward way from standard facts in linear algebra.

Lemma 2.5. *Let $A : \mathbb{R}^n \rightarrow \mathbb{R}^n$ be an affine-linear map given by $A(\vec{x}) := (W|b)\vec{x}'$ and let $\mathcal{S} = \{S_1, \dots, S_n\}$ be the associated affine solution set arrangement in \mathbb{R}^n described in equation 1. Then A is an invertible function if and only if \mathcal{S} is generic.*

Lemma 2.6. *Let $(n_0, \dots, n_m; 1)$ be a network architecture. For almost every $\mathbf{s} \in \mathbb{R}^D$, each \mathcal{S}_i in $(\mathcal{S}_0, \dots, \mathcal{S}_m)$ is a generic (co-oriented) hyperplane arrangement.*

3. POLYHEDRAL COMPLEXES

We will need some basic facts about the geometry and combinatorics of convex polytopes, polyhedral sets, and polyhedral complexes. We quickly recall relevant background and terminology, referring the interested reader to [9, 10] for a more thorough treatment.

A *polyhedral set* \mathcal{P} in \mathbb{R}^n is an intersection of finitely many closed affine half spaces $H_1^+, \dots, H_m^+ \subseteq \mathbb{R}^n$. A *convex polytope* in \mathbb{R}^n is a bounded polyhedral set. Note that a polyhedral set is an intersection of convex sets, and hence convex. Each region of a hyperplane arrangement is the interior of a polyhedral set.

A hyperplane H in \mathbb{R}^n is a *cutting hyperplane* of a polyhedral set \mathcal{P} and is said to *cut* \mathcal{P} if there exists $x_1, x_2 \in \mathcal{P}$ with $x_1 \in \mathcal{P} \cap H^+$ and $x_2 \in \mathcal{P} \cap H^-$. A hyperplane H in \mathbb{R}^n is a *supporting hyperplane* of \mathcal{P} and is said to *support* \mathcal{P} if H does not cut \mathcal{P} and $H \cap \mathcal{P} \neq \emptyset$.

For any set $S \subset \mathbb{R}^n$, the *affine hull* of S , denoted $\text{aff}(S)$, is the intersection of all affine-linear subspaces of \mathbb{R}^n containing S . The *dimension* of a polyhedral set is the dimension of its affine hull.

Let \mathcal{P} be a polyhedral set of dimension n . A subset $F \subset \mathcal{P}$ is said to be a *face* of \mathcal{P} if either $F = \emptyset$, $F = \mathcal{P}$, or $F = H \cap \mathcal{P}$ for some supporting hyperplane of \mathcal{P} . \emptyset and \mathcal{P} are called the *improper* faces of \mathcal{P} . All other faces are *proper*. A k -*face* of \mathcal{P} is a face of \mathcal{P} that has dimension k . A *facet* of \mathcal{P} is an $(n-1)$ -face of \mathcal{P} . A *vertex* of \mathcal{P} is a 0-face of \mathcal{P} .

Lemma 3.1. [9, Sec 26] *Every polyhedral set \mathcal{P} has an irredundant realization as an intersection $\mathcal{P} = H_1^+ \cap \dots \cap H_m^+$ satisfying the property that $\mathcal{P} \neq (\mathcal{P}_i := \bigcap_{j \neq i} H_j^+)$ for each $i = 1, \dots, m$. Moreover, for an irredundant realization as above, the set of facets of \mathcal{P} is precisely the set of proper faces of the form $\mathcal{P} \cap H_i$.*

A *polyhedral complex* \mathcal{C} of dimension d is a finite set of polyhedral sets of dimension k , for $0 \leq k \leq d$, called the *cells* of \mathcal{C} , such that i) If $P \in \mathcal{C}$, then every face of P is in \mathcal{C} , and ii) if $P, Q \in \mathcal{C}$, then $P \cap Q$ is a single mutual face of P and Q . The *domain* or *underlying set* $|\mathcal{C}|$ of a polyhedral complex \mathcal{C} is the union of its cells. If \mathcal{C} is a polyhedral complex embedded in \mathbb{R}^n and $|\mathcal{C}| = \mathbb{R}^n$, we call \mathcal{C} a *polyhedral decomposition* of \mathbb{R}^n .

A *polyhedral subcomplex* of \mathcal{C} is a subset $\mathcal{C}' \subseteq \mathcal{C}$ such that for every cell P in \mathcal{C}' , every face of P is also in \mathcal{C}' . The k -*skeleton* of \mathcal{C} , denoted \mathcal{C}_k , is the subcomplex of all polyhedral sets of \mathcal{C} of dimension i , where $0 \leq i \leq k$.

Any hyperplane arrangement \mathcal{A} in \mathbb{R}^n induces a polyhedral decomposition, $\mathcal{C}(\mathcal{A})$, of \mathbb{R}^n as follows. Define the n -dimensional cells of $\mathcal{C}(\mathcal{A})$ to be the closures of the regions of \mathcal{A} , and for $0 < i < n$, inductively define the i -dimensional cells of $\mathcal{C}(\mathcal{A})$ to be the facets of the $i+1$ dimensional cells. Similarly, any affine solution set arrangement $\mathcal{S} = \{S_1, \dots, S_m\}$ in \mathbb{R}^n induces a polyhedral decomposition, $\mathcal{C}(\mathcal{S})$, formed by first removing the

degenerate affine solution sets from \mathcal{S} to obtain a hyperplane arrangement, \mathcal{A} , and setting $\mathcal{C}(\mathcal{S}) := \mathcal{C}(\mathcal{A})$.

For polyhedral complexes M and R , a map $f : |M| \rightarrow |R|$ is *cellular* if for every cell $K \in M$ there exists a cell $L \in R$ with $f(K) \subseteq L$. For a polyhedral complex M embedded in \mathbb{R}^m , a map $f : |M| \rightarrow \mathbb{R}^r$ is *affine-linear on cells of M* , if for each cell $K \in M$, the restriction of f to $|K|$ is affine-linear.³

A polyhedral complex M' is said to be a *subdivision* of a polyhedral complex M if $|M| = |M'|$ and each cell of M' is contained in a cell of M .

Of particular interest here is the *level set complex* associated to a pair of polyhedral complexes M and R , with R embedded in \mathbb{R}^r , and a map $f : |M| \rightarrow \mathbb{R}^r$ affine-linear on cells of M . The level set complex, which is a polyhedral complex by [10, Lemma 2.5], is the complex

$$M_{\in R} := \{S \cap f^{-1}(Y) \mid S \in M, Y \in R\}.$$

In the present work, we focus primarily on maps $\mathbb{R}^n \rightarrow \mathbb{R}$, where the single 0-cell of the polyhedral complex $R \subseteq \mathbb{R}$ is a threshold $t \in \mathbb{R}$ and the two 1-cells are the unbounded intervals $(-\infty, t]$ and $[t, \infty)$.

4. TRANSVERSALITY

In this subsection, we recall classical transversality results and state and extend them to situations of relevance to us.

4.1. Classical transversality results. We follow the terminology and notation of [11].

Denote the tangent space of a smooth manifold X at a point $x \in X$ by $T_x X$. Recall that for a smooth map $f : X \rightarrow Y$ of manifolds with $f(x) = y$, the derivative df_x is a linear map between tangent spaces, $df_x : T_x X \rightarrow T_y Y$, and the image $df_x(T_x X)$ is a linear subspace of $T_y Y$. If U and V are two subspaces of a linear space W , then their sum, $U + V$, is the subspace $\{u + v : u \in U, v \in V\}$.

In Definition 4.1 and Theorems 7 and 8, assume X to be a smooth manifold with or without boundary, Y and Z to be smooth manifolds without boundary, Z a smoothly embedded submanifold of Y , and $f : X \rightarrow Y$ a smooth map.

Definition 4.1. *We say that f is transverse to Z and write $f \pitchfork Z$ if*

$$(2) \quad df_p(T_p X) + T_{f(p)} Z = T_{f(p)} Y$$

for all $p \in f^{-1}(Z)$.

Note that Definition 4.1 allows for the possibility that X is a manifold of dimension 0, i.e. consists of—without loss of generality—a single point p . In this case $T_p\{p\} = \{0\}$ and so $df_p(T_p\{p\}) = \{0\}$, so condition (2) reduces to the condition that if $f(p) \in Z$, then Z and Y must agree in a neighborhood of $f(p)$. Note also that if the image $f(X)$ does not intersect Z , then condition (2) is vacuously true, hence $f \pitchfork Z$.

³Note that [10], following the classical texts on PL geometry, uses “linear on cells” rather than “affine-linear on cells,” since PL charts need not specify an origin. Since our polyhedral complexes are canonically embedded in \mathbb{R}^n , we use “affine-linear” throughout.

Theorem 7 (Map Transversality Theorem). [11, p. 28] *If f is transverse to Z , then $f^{-1}(Z)$ is an embedded submanifold of X . Furthermore, the codimension of $f^{-1}(Z)$ in X equals the codimension of Z in Y .*

The Map Transversality Theorem uses the standard convention that the empty set can assume any dimension. If $f^{-1}(Z) = \emptyset$, one considers the codimension of $f^{-1}(Z)$ in X to be the codimension of Z in Y .

Theorem 8 (Parametric Transversality Theorem). [11, p. 68] *Let S be a smooth manifold and let $F : X \times S \rightarrow Y$ be a smooth map. If F is transverse to Z , then for (Lebesgue) almost every $s \in S$ the restriction map $F_s : X \rightarrow Y$ given by $F_s(x) = F(x, s)$ is transverse to Z .*

We wish to apply Theorem 8 to the parametrized family of neural networks of a fixed architecture (Definition 2.3), but this family does not satisfy the smoothness requirements, so we develop the necessary non-smooth analogues in §4.2.

4.2. Extensions of the classical transversality results to maps on polyhedral complexes that are smooth on cells.

We introduce a polyhedral analogue of Definition 4.1:

Definition 4.2. *Let X be a polyhedral complex of dimension d in \mathbb{R}^n , let $f : |X| \rightarrow \mathbb{R}^r$ be a map which is smooth on all cells of X and let Z be a smoothly embedded submanifold (without boundary) of \mathbb{R}^r . We say that f is transverse on cells to Z and write $f \pitchfork_c Z$ if:*

- (i) *the restriction of f to the interior, $\text{int}(C)$, of every k -cell C of X is transverse to Z (in the sense of Definition 4.1) when $1 \leq k \leq d$, and*
- (ii) *the restriction of f to every 0-cell of X is transverse to Z .*

Note that a function defined on a 0-cell is considered to be smooth. We will be particularly interested in the case in which $r = 1$ and $Z = \{t\}$ is a threshold in \mathbb{R} .

Corollary 4.3. *Let X be a polyhedral complex of dimension d in \mathbb{R}^n . Let $f : |X| \rightarrow \mathbb{R}^r$ be a map which is smooth on cells of X and let Z be a smoothly embedded submanifold of \mathbb{R}^r for which $f \pitchfork_c Z$. Then we have:*

- *For every k -cell $C \in X$, where $1 \leq k \leq d$, $f^{-1}(Z) \cap \text{int}(C)$ is a (possibly empty) smoothly embedded submanifold of $\text{int}(C)$. Furthermore, the codimension of $f^{-1}(Z) \cap \text{int}(C)$ in $\text{int}(C)$ equals the codimension of Z in \mathbb{R}^r .*
- *If $\dim(Z) < r$, then for every 0-cell $C \in X$ (vertex), $f(C) \notin Z$.*

Proof. This follows immediately from Theorem 7, since the interior of any polyhedral set of dimension $k \in \mathbb{N}$ is a nonempty smooth manifold. Note also that condition (ii) of Definition 4.2 implies that if f is transverse on cells to Z and there exists a vertex v of X such that $f(v) \in Z$, then Z must have the full dimension r . Thus, if $\dim(Z) < r$, f being transverse on cells to Z implies no vertex of X is sent by f to Z . \square

We will need the following version of Theorem 8 for families of maps that are smooth on cells of a polyhedral complex.

Proposition 4.4. *Let X be a polyhedral complex in \mathbb{R}^n , S a smooth manifold without boundary, and $Z \subseteq \mathbb{R}^r$ a smoothly embedded submanifold without boundary. Suppose that for each cell $C \in X$, the restricted map $F|_{C \times S} : C \times S \rightarrow \mathbb{R}^r$ is smooth and the further restricted map $F|_{C' \times S} : C' \times S \rightarrow \mathbb{R}^r$, where*

$$C' = \begin{cases} \text{int}(C) & \text{if } C \text{ is of dimension } \geq 1, \\ C & \text{if } C \text{ is of dimension } 0, \end{cases}$$

is transverse to Z . Then for (Lebesgue) almost every $s \in S$, the map

$$f_s : |X| \rightarrow \mathbb{R}^r$$

given by $f_s(x) = F(x, s)$ is transverse on cells to Z .

Proof. For each cell $C \in X$, the Parametric Transversality Theorem implies that there exists a null set $S_C \subset S$ such that $f_s|_{C'}$ is transverse to Z for every $s \in S \setminus S_C$. Let $S_X = \bigcup_{C \in X} S_C$; as a finite union of null sets, S_X is a null set. Then for every $C \in X$ and $s \in S \setminus S_X$, we have that $f_s|_{C'} : C' \rightarrow \mathbb{R}^r$ is transverse to Z . Hence f_s is transverse on cells to Z for all $s \in S \setminus S_X$. \square

5. MAPS ON POLYHEDRAL COMPLEXES AND TRANSVERSAL THRESHOLDS

We now turn to applying the transversality statements developed in the previous section to ReLU neural network maps.

Definition 5.1. *Let M be a polyhedral complex embedded in \mathbb{R}^{n_0} , $n_0 \in \mathbb{N}$, and let $F : |M| \rightarrow \mathbb{R}$ be a map that is smooth on cells. A threshold $t \in \mathbb{R}$ is said to be transversal for F and M if F is transverse on cells (Definition 4.2) to the submanifold $\{t\} \subseteq \mathbb{R}$. In this case, we write $F \pitchfork_c \{t\}$.*

Although Section 4.2 and Definition 5.1 require only that the map F be smooth on cells, from this point onwards we restrict to the case that F is affine-linear on cells, since this is the setting relevant for understanding neural network maps. For the remainder of this section, let M be a polyhedral complex embedded in \mathbb{R}^{n_0} , $n_0 \in \mathbb{N}$, and let $F : |M| \rightarrow \mathbb{R}$ be a map that is affine-linear on cells.

Definition 5.2. *A point $x \in M$ is said to have a F -nonconstant cellular neighborhood in M if F is nonconstant on each cell of M containing x .*

Note that each vertex of M is itself a cell on which F is necessarily constant; hence, no vertex of M can be said to have a F -nonconstant cellular neighborhood.

Lemma 5.3. *A threshold $t \in \mathbb{R}$ is transversal for F and M if and only if each point $p \in F^{-1}(\{t\})$ has a F -nonconstant cellular neighborhood in M .*

Proof. The threshold t is transversal for F and M if and only if for any k -cell $C \in X$ with $k \geq 1$ the restriction of f to $\text{int}(C)$ is transverse to $\{t\}$, and the restriction of f to any 0-cell X is transverse to $\{t\}$. This is equivalent to the statement that if $p \in F^{-1}(\{t\})$ and p is in a cell $C \in M$, then

$$df_p(T_p C) + T_{f(p)}\{t\} = T_{f(p)}\mathbb{R}.$$

Since $T_{f(p)}\{t\} = \{0\}$, this equality holds if and only if $df_p(T_p C) = T_{f(p)}\mathbb{R} \cong \mathbb{R}$ for every cell C containing p . By Corollary 4.3, this is equivalent to p having a nonconstant cellular neighborhood, as desired. \square

Lemma 5.4. *Let $t \in \mathbb{R}$ be a transversal threshold for F and M . Then for every cell $C \in M$, $F^{-1}(\{t\}) \cap C$ is either empty or $\text{aff}(F^{-1}(\{t\}) \cap C)$ is a hyperplane in $\text{aff}(C)$. Moreover, whenever $F^{-1}(\{t\}) \cap C$ is nonempty, the hyperplane $\text{aff}(F^{-1}(\{t\}) \cap C)$ cuts C .*

Proof. The statement that $F^{-1}(\{t\}) \cap C$ is a submanifold of codimension 1 in C is from the Map Transversality Theorem (Theorem 7); its affine hull is a hyperplane because F is affine-linear. Let $H = \text{aff}(F^{-1}(\{t\}) \cap C) \neq \emptyset$. If H were a supporting hyperplane of C , then $H \cap C$ would be a non-empty lower-dimensional face of C , all of whose points map to t . Applying Lemma 5.3, this would contradict the assumption that t is a transversal threshold. Hence, H cuts C whenever $F^{-1}(\{t\}) \cap C \neq \emptyset$. \square

Lemma 5.5. *All but finitely many thresholds $t \in \mathbb{R}$ are transversal for F and M .*

Proof. The polyhedral complex M is, by definition, finite. Hence there are only finitely many cells on which F is constant. But Lemma 5.3 tells us that the images of the constant cells are the only nontransversal thresholds for F and M . \square

6. BENT HYPERPLANE ARRANGEMENTS AND CANONICAL POLYHEDRAL COMPLEXES

The following notion was introduced in [14].

Definition 6.1. [14, Eqn. (2), Lem. 4] *Let*

$$\mathbb{R}^{n_0} \xrightarrow{F_1=\sigma \circ A_1} \mathbb{R}^{n_1} \xrightarrow{F_2=\sigma \circ A_2} \dots \xrightarrow{F_m=\sigma \circ A_m} \mathbb{R}^{n_m} \xrightarrow{G=A_{m+1}} \mathbb{R}^1.$$

be a ReLU neural network and let $\mathcal{A}^{(k)} = \{H_1^{(k)}, \dots, H_{n_{i_k}}^{(k)}\}$ denote the hyperplane arrangement in $\mathbb{R}^{n_{k-1}}$ associated to the layer map F_k .

A bent hyperplane associated to the k^{th} layer of F , for $k \in \{2, \dots, m\}$, is the preimage in \mathbb{R}^{n_0} of any hyperplane $H_i^{(k)} \subseteq \mathbb{R}^{n_{k-1}}$ associated to the k^{th} layer map:

$$(F_{k-1} \circ \dots \circ F_1)^{-1} \left(H_i^{(k)} \right).$$

Although they do not bend, for the sake of consistency we will refer to the hyperplanes $H_i^{(1)}$ in $\mathcal{A}^{(1)} \subseteq \mathbb{R}^{n_0}$ as the bent hyperplanes associated to the 1st layer of F .

We will denote by $\mathcal{B}_F^{(k)}$ the union of the bent hyperplanes associated to the k^{th} layer of F and refer to this union as the *bent hyperplane arrangement associated to the k^{th} layer of F* .

We will denote by $\mathcal{B}_F := \bigcup_{k=1}^m \mathcal{B}_F^{(k)}$ the union of all bent hyperplanes from all m layers of F and refer to this union as the *bent hyperplane arrangement associated to F* .

Remark 6.2. Since there is no activation function on the final layer map to induce any additional loci of non-differentiability, we do not include its associated bent hyperplane in \mathcal{B}_F .

It is immediate that F is smooth on the complement of \mathcal{B}_F .

Definition 6.3. [14, Def 1, Lem 2] Let $F : \mathbb{R}^{n_0} \rightarrow \mathbb{R}$ be a ReLU neural network map. An activation region of F is a connected component of the complement of the bent hyperplane arrangement associated to F , i.e. a connected component of $\mathbb{R}^{n_0} \setminus \mathcal{B}_F$.

Remark 6.4. Note that for $k \geq 2$ it is possible for $\mathcal{B}_F^{(k)}$ to have codimension 0, not 1, in \mathbb{R}^{n_0} . As a simple example of this phenomenon, consider a two-layer ReLU neural network

$$F : \mathbb{R}^2 \xrightarrow{F_1} \mathbb{R}^2 \xrightarrow{F_2} \mathbb{R} \xrightarrow{G} \mathbb{R},$$

where $\mathcal{A}^{(1)}$ is the standard coordinate hyperplane arrangement, and $\mathcal{A}^{(2)} = \{H^{(2)}\}$, where $H^{(2)}$ is any line through the origin with negative slope. Then

$$\mathcal{B}_F^{(1)} = \{H_1^{st}, H_2^{st}\},$$

the standard co-oriented coordinate axes, and

$$\mathcal{B}_F^{(2)} = \{(x, y) \in \mathbb{R}^2 \mid x, y \leq 0, \}$$

the closed non-positive orthant. In particular, the bent hyperplane arrangement is codimension 0, not 1, and hence the closure of the activation regions (Definition 6.3) is a proper subset of \mathbb{R}^n .

This phenomenon arises when a map fails to be transversal to a threshold, an observation that motivates Definition 8.1 and Theorem 2. Note that it is also a measure zero phenomenon. See Theorem 3.

We now define a canonical polyhedral decomposition of the domain of a ReLU neural network. In the transversal case, we explicitly relate this decomposition to the bent hyperplane arrangements and activation regions in Theorems 1 and 2.

Definition 6.5. Let

$$F : \mathbb{R}^{n_0} \xrightarrow{F_1} \mathbb{R}^{n_1} \xrightarrow{F_2} \dots \xrightarrow{F_m} \mathbb{R}^{n_m} \xrightarrow{G} \mathbb{R}$$

be a ReLU neural network. For $i \in \{1, \dots, m\}$, denote by $R^{(i)}$ the polyhedral complex in $\mathbb{R}^{n_{i-1}}$ induced by the hyperplane arrangement associated to the i^{th} layer map F_i . Inductively define polyhedral complexes, $\mathcal{C}(F_i \circ \dots \circ F_1)$, in \mathbb{R}^{n_0} as follows: Set

- $\mathcal{C}(F_1) := R^{(1)}$, and
- $\mathcal{C}(F_i \circ \dots \circ F_1) := \mathcal{C}(F_{i-1} \circ \dots \circ F_1)_{\in R^{(i)}}$ for $i = 2, \dots, m$.

The canonical polyhedral complex associated to F is $\mathcal{C}_F := \mathcal{C}(F_m \circ \dots \circ F_1)$.

Theorem 1. Let

$$F : \mathbb{R}^{n_0} \xrightarrow{F_1} \mathbb{R}^{n_1} \xrightarrow{F_2} \dots \xrightarrow{F_m} \mathbb{R}^{n_m} \xrightarrow{G} \mathbb{R}$$

be a ReLU neural network. For each i , $\mathcal{C}(F_i \circ \dots \circ F_1)$ is a polyhedral decomposition of \mathbb{R}^{n_0} satisfying

- (i) $F_i \circ \dots \circ F_1$ is affine-linear on the cells of $\mathcal{C}(F_i \circ \dots \circ F_1)$,
- (ii) $\bigcup_{k=1}^i \mathcal{B}_F^{(k)}$ is the domain of a polyhedral subcomplex of $\mathcal{C}(F_i \circ \dots \circ F_1)$.

Proof. For each $i = 1, \dots, m$, denote by $\mathcal{A}^{(i)}$ the hyperplane arrangement associated to the layer map $F_i : \mathbb{R}^{n_{i-1}} \rightarrow \mathbb{R}^{n_i}$ and denote by $R^{(i)}$ the induced polyhedral decomposition of $\mathbb{R}^{n_{i-1}}$. We proceed by induction on i .

For $i = 1$, it is immediate that $\mathcal{B}_F^{(1)} = \mathcal{A}^{(1)}$ forms the $(n_0 - 1)$ -skeleton of $R^{(1)}$ and F_1 is affine-linear on cells of $\mathcal{C}(F_1) = R^{(1)}$.

Now consider $i > 1$ and assume the statement holds for $i - 1$. Since level set complexes are polyhedral complexes, condition (i) of the inductive hypothesis implies $\mathcal{C}(F_i \circ \dots \circ F_1)$ is a polyhedral complex.

By condition (i) of the inductive hypothesis, each cell in $\mathcal{C}(F_i \circ \dots \circ F_1)$ is the intersection of a cell in $\mathcal{C}(F_{i-1} \circ \dots \circ F_1)$ with the preimage of a cell in $R^{(i)}$. The map $F_{i-1} \circ \dots \circ F_1$ is affine-linear on each such intersection by assumption. The layer map $F_i : \mathbb{R}^{n_{i-1}} \rightarrow \mathbb{R}^{n_i}$ is affine-linear on cells of $R^{(i)}$. Condition (i) follows.

By condition (ii) of the inductive hypothesis, $\bigcup_{k=1}^{i-1} \mathcal{B}_F^{(k)}$ is the domain of a polyhedral subcomplex of $\mathcal{C}(F_{i-1} \circ \dots \circ F_1)$. By definition, $\mathcal{C}(F_i \circ \dots \circ F_1)$ is a subdivision of $\mathcal{C}(F_{i-1} \circ \dots \circ F_1)$, so $\bigcup_{k=1}^{i-1} \mathcal{B}_F^{(k)}$ is the domain of a polyhedral subcomplex of $\mathcal{C}(F_i \circ \dots \circ F_1)$. Let $R_{(n_{i-1}-1)}^{(i)}$ denote the $(n_{i-1} - 1)$ -skeleton of $R^{(i)}$. Noting that the domain of $R_{(n_{i-1}-1)}^{(i)}$ is the union of the hyperplanes in $\mathcal{A}^{(i)}$, we have

$$|\mathcal{C}(F_{i-1} \circ \dots \circ F_1)|_{\in R_{(n_{i-1}-1)}^{(i)}} = \mathcal{B}_F^{(i)}.$$

Since the union of two subcomplexes of a polyhedral complex is a subcomplex, $\bigcup_{k=1}^i \mathcal{B}_F^{(k)}$ is a polyhedral subcomplex of $\mathcal{C}(F_i \circ \dots \circ F_1)$, implying condition (ii). \square

The following definition is important for the constructions in Section 8:

Definition 6.6. A threshold $t \in \mathbb{R}$ is a transversal threshold for a neural network $F : \mathbb{R}^{n_0} \rightarrow \mathbb{R}$ if t is a transversal threshold for F and its canonical polyhedral complex $\mathcal{C}(F)$.

7. PIECEWISE SMOOTHNESS OF THE PARAMETRIZED FAMILY OF NEURAL NETWORKS

Throughout this section, consider any fixed architecture $(n_0, \dots, n_m; 1)$ and let $\mathcal{F} : \mathbb{R}^{n_0} \times \mathbb{R}^D \rightarrow \mathbb{R}$ be the parametrized family of neural networks of this architecture (Definition 2.3).

Lemma 7.1. There exists a finite set E of polynomials in the variables $x_1, \dots, x_{n_0}, s_1, \dots, s_D$ such that the following hold:

- (i) If T is a term of a polynomial in E , then T has the form $xs_1^{\tau_1} \dots s_D^{\tau_D}$ for some $x \in \{x_1, \dots, x_{n_0}, 1\}$ and $(\tau_1, \dots, \tau_D) \in \{0, 1\}^D$.
- (ii) \mathcal{F} is smooth on the complement of the set Z defined by

$$Z := \{(x, s) \in \mathbb{R}^{n_0} \times \mathbb{R}^D : f_i((x, s)) = 0 \text{ for some } f_i \in E\}.$$

Proof. The idea is to let E be the set of all possible “inputs” of any ReLU in the expression defining \mathcal{F} . Rather than presenting a formal proof, we give an illustrative example that demonstrates all the key ideas.

We consider the network architecture $(1, 2, 1; 1)$. The associated parametrized family is the map $\mathcal{F} : \mathbb{R} \times \mathbb{R}^9 \rightarrow \mathbb{R}$ given by

$$(x, (a, b, c, d, e, f, g, h, i)) \mapsto h \cdot \text{ReLU}(e \cdot \text{ReLU}(ax+b) + f \cdot \text{ReLU}(cx+d) + g) + i$$

Each of the three ReLU's in this expression acts as either the identity or 0, depending on the sign of its argument. Let E be the set of all possible expressions that are inputs of a ReLU in the expression for \mathcal{F} , allowing for the possibility that each nested ReLU could be either 0 or the identity. That is,

$$E = \{ax + b, cx + d, e(ax + b) + f(cx + d) + g, e(ax + b) + g, f(cx + d) + g, g\}.$$

Let $Z \subset \mathbb{R} \times \mathbb{R}^9$ be the set of points (x, s) where at least one function in E evaluates to 0.

For any fixed input (x, s) , $\mathcal{F}((x, s))$ is given by one of 2^3 possible (not necessarily distinct) formulas (which correspond to each of the 3 ReLU's being in one of two possible "states"). Let H be the set of 2^3 (not necessarily distinct) functions formed by replacing each ReLU with either 0 or the identity. Which of these expressions represents \mathcal{F} locally only changes at points where the argument of a ReLU in the expression for \mathcal{F} – i.e. a polynomial in E – changes sign. Now, since all the functions in E are continuous and Z is their set of zeros, for any point $(x, s) \in (\mathbb{R} \times \mathbb{R}^9) \setminus Z$ there exists a neighborhood U of (x, s) on which the sign of each function in E is constant. Consequently, there is a fixed function $f \in H$ such that \mathcal{F} agrees with f on U . Since all functions in H are smooth, it follows that the restriction of \mathcal{F} to $(\mathbb{R} \times \mathbb{R}^9) \setminus Z$ is smooth. \square

Theorem 4. *There exists an algebraic set $Z \subset \mathbb{R}^{n_0} \times \mathbb{R}^D$ such that*

- (i) \mathcal{F} is smooth on the complement of Z ,
- (ii) Z is the vanishing set of a polynomial, and hence is a closed, nowhere dense subset with Lebesgue measure 0, and
- (iii) the complement of Z consists of finitely many connected components.

Proof. Let E be the set of polynomials constructed in Lemma 7.1, define the polynomial $F := \prod_{f_i \in E} f_i$, and observe that the set $Z \subset \mathbb{R}^{n_0+D}$ from Lemma 7.1 is the vanishing set for the polynomial F . Items (i) and (ii) follow immediately. To see (iii), consider the polynomial $G : \mathbb{R}^{D+n_0} \times \mathbb{R} \rightarrow \mathbb{R}$ defined by $G(x, z) = zF(x) - 1$, and let $Z_G \subset \mathbb{R}^{D+n_0} \times \mathbb{R}^1$ be the vanishing set for G . As the vanishing set of a real polynomial, Z_G is a real algebraic variety; hence Z_G has finitely many connected components (by e.g. [26, Theorem 3] or [25, Lemma 2.5]). Note that $G(x, z) = 0$ if and only if $F(x) \neq 0$, $z = 1/F(x)$. Hence Z_G is homeomorphic (via $(x, z) \mapsto x$) to $\mathbb{R}^{D+n_0} \setminus Z$; consequently $\mathbb{R}^{D+n_0} \setminus Z$ also has finitely many connected components. \square

Proposition 7.2. *Let Z be the algebraic set constructed in the proof of Lemma 7.1. For each $s \in \mathbb{R}^D$, denote by \mathcal{F}_s the neural network map defined by $\mathcal{F}_s(x) = \mathcal{F}(x, s)$. Then $\bigcup_{(x, s) \in Z} (\mathcal{B}_{\mathcal{F}_s} \times \{s\}) \subseteq Z$.*

Proof. Fix a parameter $s \in \mathbb{R}^D$. The bent hyperplane arrangement $\mathcal{B}_{\mathcal{F}_s}$ is the union of the preimages in \mathbb{R}^{n_0} of the hyperplanes in the hyperplane arrangements associate to the layer maps of \mathcal{F}_s . Each such hyperplane is the vanishing set of the argument of a ReLU in the expression for \mathcal{F}_s . Thus the polynomial that defines such a hyperplane is obtained by substituting the values for s into one of the polynomials in the set E constructed in the proof of Lemma 7.1, implying $F(x, s) = 0$, where F is the product of the polynomials in E . Hence, any point x in the preimage in \mathbb{R}^{n_0} of such a hyperplane satisfies $(x, s) \in Z$. \square

8. TRANSVERSAL NEURAL NETWORKS

In what follows, let $\pi_j : \mathbb{R}^m \rightarrow \mathbb{R}$ denote the projection onto the j^{th} coordinate. For any neural network map

$$\mathbb{R}^{n_0} \xrightarrow{F_1=\sigma \circ A_1} \mathbb{R}^{n_1} \xrightarrow{F_2=\sigma \circ A_2} \dots \xrightarrow{F_m=\sigma \circ A_m} \mathbb{R}^{n_m} \xrightarrow{G=A_{m+1}} \mathbb{R}^1,$$

$i \in \{1, \dots, m\}$ and $j \in \{1, \dots, n_i\}$, the *node map*, $F_{i,j}$, is the map

$$\pi_j \circ A_i \circ F_{i-1} \circ \dots \circ F_1 : \mathbb{R}^{n_0} \rightarrow \mathbb{R}.$$

Definition 8.1. A ReLU neural network

$$F : \mathbb{R}^{n_0} \xrightarrow{F_1} \mathbb{R}^{n_1} \xrightarrow{F_2} \dots \xrightarrow{F_m} \mathbb{R}^{n_m} \xrightarrow{G} \mathbb{R}$$

is said to be *transversal* if, for each $i \in \{1, \dots, m\}$ and each $j \in \{1, \dots, n_i\}$, $t = 0$ is a transversal threshold (Definition 5.1) for the node map

$$F_{i,j} : \mathbb{R}^{n_0} \rightarrow \mathbb{R}.$$

Remark 8.2. The descriptors *generic* and *transversal*, when applied to ReLU neural networks, are similar but complementary concepts.

A ReLU neural network is *generic* if each solution set arrangement for each layer map is generic. It is not immediate, yet it is true, that if a solution set arrangement is generic then each solution set in the arrangement intersects each *intersection* of solution sets in that layer transversely.

In contrast, if a ReLU neural network is *transversal* then it follows from the definitions that each bent hyperplane intersects the bent hyperplanes from all previous layers transversely.

Put simply, when applied to ReLU neural networks, the term *generic* describes intersections of cells associated to a single layer map, and the term *transversal* describes intersections of cells associated to different layer maps.

Theorem 2. If a ReLU neural network $F : \mathbb{R}^{n_0} \rightarrow \mathbb{R}$ is transversal, then the bent hyperplane arrangement \mathcal{B}_F is the domain of the $(n_0 - 1)$ -skeleton of the canonical polyhedral complex $\mathcal{C}(F)$, and the closures of the activation regions of F are the n_0 -cells of $\mathcal{C}(F)$.

Proof. We proceed by induction on m , the number of hidden layers of F . The base case $m = 1$ is immediate. Now consider any fixed value of $m > 1$ and assume the result holds for all smaller values of m . In particular, for each node map $F_{m,j}$ the bent hyperplane arrangement $\mathcal{B}_{F_{m,j}} = \bigcup_{i=1}^{m-1} \mathcal{B}_{F_{m,j}}^{(i)}$

is the domain of the $(n_0 - 1)$ -skeleton of $\mathcal{C}(F_{m-1} \circ \dots \circ F_1)$. But for each $i \in \{1, \dots, m-1\}$ and $j \in \{1, \dots, n_m\}$, we have $\mathcal{B}_{F_{m,j}}^{(i)} = \mathcal{B}_F^{(i)}$, so the bent hyperplane arrangement,

$$\mathcal{B}'_F := \bigcup_{i=1}^{m-1} \mathcal{B}_F^{(i)},$$

for the first $m-1$ layers is the $(n_0 - 1)$ -skeleton of the polyhedral complex $\mathcal{C}(F_{m-1} \circ \dots \circ F_1)$.

To see that \mathcal{B}_F is contained in the $(n_0 - 1)$ -skeleton of $\mathcal{C}(F_m \circ \dots \circ F_1)$, we begin by noting that \mathcal{B}'_F is contained in the $(n_0 - 1)$ -skeleton of $\mathcal{C}(F_m \circ \dots \circ F_1)$ since $\mathcal{C}(F_m \circ \dots \circ F_1)$ is a subdivision of $\mathcal{C}(F_{m-1} \circ \dots \circ F_1)$. Moreover, by definition

$$\mathcal{B}_F = \mathcal{B}'_F \cup \bigcup_{j=1}^{n_m} F_{m,j}^{-1}(\{0\}).$$

It therefore suffices to show that $\bigcup_{j=1}^{n_m} F_{m,j}^{-1}(\{0\})$ is contained in the $(n_0 - 1)$ -skeleton of $\mathcal{C}(F_m \circ \dots \circ F_1)$.

But since 0 is a transversal threshold for each $F_{m,j} : \mathbb{R}^{n_0} \rightarrow \mathbb{R}$, this follows from Corollary 4.3. Explicitly, for every cell $C \in \mathcal{C}(F_{m-1} \circ \dots \circ F_1)$, Theorem 7 tells us that $C \cap F_{m,j}^{-1}(\{0\})$ is codimension 1 in C . Since $\mathcal{C}(F_{m-1} \circ \dots \circ F_1)$ has dimension n_0 , it follows that any new cell in $\mathcal{B}_F \setminus \mathcal{B}'_F$ has dimension $\leq (n_0 - 1)$, as desired.

To see that $\mathcal{C}(F_m \circ \dots \circ F_1)$ is contained in \mathcal{B}_F we will show that any k -cell C in $\mathcal{C}(F_m \circ \dots \circ F_1)_{n_0-1}$ is also in \mathcal{B}_F . Since $\mathcal{C}(F_m \circ \dots \circ F_1)$ is, by definition, a subdivision of $\mathcal{C}(F_{m-1} \circ \dots \circ F_1)$, the cell C is contained in a cell C' of $\mathcal{C}(F_{m-1} \circ \dots \circ F_1)$. Assume WLOG that C' has minimal dimension among all cells in $\mathcal{C}(F_{m-1} \circ \dots \circ F_1)$ containing C , and let k' be the dimension of C' . If $k' \leq n_0 - 1$, then $C \subseteq C' \subseteq \mathcal{C}(F_{m-1} \circ \dots \circ F_1)_{n_0-1}$, and the inductive hypothesis tells us $C \subseteq \mathcal{B}'_F$, as desired.

So we may assume that C' has dimension n_0 . Therefore the construction described in the proof of Theorem 1 and the fact that C has dimension $\leq n_0 - 1$ (see also Lemma 5.4) tells us that C is equal to the intersection of C' with $F_{m,j}^{-1}(\{0\})$ for some node j in the m th layer map. It follows that $C \subseteq \mathcal{B}_F$. We conclude that $\mathcal{B}_F = \mathcal{C}(F_m \circ \dots \circ F_1)_{n_0-1}$ as desired. \square

Theorem 3. *For any given architecture $(n_0, \dots, n_m; 1)$ of feedforward ReLU neural network, almost every (with respect to Lebesgue measure on \mathbb{R}^D) choice of parameters yields a transversal ReLU neural network.*

Proof. We proceed by induction on m , the number of hidden layers. In the base case $m = 0$, for any network map F of this architecture, the unique node map $F_{1,1}$ is the affine-linear map $G : \mathbb{R}^{n_0} \rightarrow \mathbb{R}$, and every such map is transversal.

Now consider a fixed value of $m \geq 1$, and assume the result holds for smaller values of m . For any neural network F of architecture $(n_0, \dots, n_m; 1)$ and any $j \in \{1, \dots, n_m\}$, the node map $F_{m,j}$ is a neural network of architecture $(n_0, \dots, n_{m-1}; 1)$. Note that such a neural network F is transversal if and only if the node map $F_{m,1}$ is a transversal neural network and for every $j \in \{1, \dots, n_m\}$ the node map $F_{m,j}$ has $t = 0$ as a transversal threshold.

For any $n \in \mathbb{N}$, let λ_n denote Lebesgue measure on \mathbb{R}^n . By the inductive hypothesis, there exists a $\lambda_{D(n_0, \dots, n_{m-1}; 1)}$ -null set $N_1 \subset \mathbb{R}^{D(n_0, \dots, n_{m-1}; 1)}$ such that for every parameter in $\mathbb{R}^{D(n_0, \dots, n_{m-1}; 1)} \setminus N_1$, the associated neural network map (which we will use as the node map $F_{m,1}$) is transversal. Note, furthermore, that transversality does not depend on the final affine-linear map to \mathbb{R}^1 , so the set N_1 has a product structure in which each coordinate corresponding to this final affine map corresponds to a factor of \mathbb{R} . For any parameter $s \in \mathbb{R}^{D(n_0, \dots, n_{m-1}; 1)}$, denote by \tilde{s} the projection of s onto the coordinates of s that determine the non-affine-linear layer of the associated network, i.e. \tilde{s} is the first $k := D(n_0, \dots, n_{m-1}; 1) - (n_{m-1} + 1)$ coordinates of s .

Set $\delta := D(n_0, \dots, n_m; 1) - k$. We will show that for each parameter $s \in \mathbb{R}^{D(n_0, \dots, n_{m-1}; 1)} \setminus N_1$, there exists a set $Y_s \subset \mathbb{R}^\delta$ such that $\lambda_\delta(Y_s) = 0$ and for every $w \in \mathbb{R}^\delta \setminus Y_s$, the neural network F associated to the parameter $(\tilde{s}, w) \in \mathbb{R}^{D(n_0, \dots, n_m; 1)}$ is such that for every $j \in \{1, \dots, n_m\}$, 0 is a transversal threshold of the node map $F_{m,j}$. Assuming such sets Y_s are defined, set

$$N_2 := \left\{ (\tilde{s}, x) : s \in \mathbb{R}^{D(n_0, \dots, n_{m-1}; 1)} \setminus N_1, x \in Y_s \right\}.$$

Tonelli's Theorem will then imply that $\lambda_{D(n_0, \dots, n_m; 1)}(N_2) = 0$. Then, for every parameter in $\mathbb{R}^{D(n_0, \dots, n_m; 1)} \setminus N_2$, the associated neural network is transversal.

So fix $s \in \mathbb{R}^{D(n_0, \dots, n_{m-1}; 1)} \setminus N_1$ and let $\mathcal{C}(F_{m-1} \circ \dots \circ F_1)$ be the associated canonical polyhedral complex. For each node index $j \in \{1, \dots, n_m\}$, consider the parametrized family

$$(\mathcal{F}_{m,j})_s : \mathbb{R}^{n_0} \times \mathbb{R}^{n_{m-1}+1} \rightarrow \mathbb{R}$$

describing the action on \mathbb{R}^{n_0} of neural networks of architecture $(n_0, \dots, n_{m-1}; 1)$ whose non-affine-linear layer maps are parameterized by \tilde{s} and whose final affine-linear map to \mathbb{R} is parameterized by the point in $\mathbb{R}^{n_{m-1}+1}$. By construction, for any choice of $w \in \mathbb{R}^{n_{m-1}+1}$, the function $(\mathcal{F}_{m,j})_s|_w : \mathbb{R}^{n_0} \rightarrow \mathbb{R}$ is affine-linear on the cells of $\mathcal{C}(F_{m-1} \circ \dots \circ F_1)$. Therefore, by Proposition 4.4, to prove that $(\mathcal{F}_{m,j})_s|_w$ is transverse to $\{0\} \subset \mathbb{R}$ for $\lambda_{n_{m-1}+1}$ -almost every $w \in \mathbb{R}^{n_{m-1}+1}$, it suffices to show that for every cell $C \in \mathcal{C}(F_{m-1} \circ \dots \circ F_1)$, the parametrized map

$$(\mathcal{F}_{m,j})_s|_{C' \times \mathbb{R}^{n_{m-1}+1}} : C' \times \mathbb{R}^{n_{m-1}+1} \rightarrow \mathbb{R}$$

is transverse to $\{0\}$, where $C' = C$ if $\dim(C) = 0$ and $C' = \text{int}(C)$ otherwise. For any such C , to show that $(\mathcal{F}_{m,j})_s|_{C' \times \mathbb{R}^{n_{m-1}+1}}$ is transverse to $\{0\} \subseteq \mathbb{R}$, it suffices to show that $(\mathcal{F}_{m,j})_s|_{C' \times \mathbb{R}^{n_{m-1}+1}}$ is surjective, since the whole space, \mathbb{R} , is clearly transverse to any embedded submanifold. But this is clearly true, because for any non-empty cell, $C \in \mathcal{C}(F_{m-1} \circ \dots \circ F_1)$, there exists a point $p \in F_{m-1} \circ \dots \circ F_1(C) \subset \mathbb{R}^{n_{m-1}}$, and it is clear that there exists *some* affine-linear transformation $\mathbb{R}^{n_{m-1}} \rightarrow \mathbb{R}$ sending p to t .

Thus, for $\lambda_{n_{m-1}+1}$ -almost every $w \in \mathbb{R}^{n_{m-1}+1}$, the node map $(\mathcal{F}_{m,j})_s|_w$ is transverse to $\{0\}$; let $Y_{s,j}$ be the $\lambda_{n_{m-1}+1}$ -null set where this fails. Set Y_s to be the set of points $w \in \mathbb{R}^\delta$ such that for some $j \in \{1, \dots, n_m\}$, the

projection of (\tilde{s}, w) to its $(n_{m-1} + 1)$ coordinates representing the affine-linear layer map of the node map $F_{m,j}$ is in $Y_{s,j}$. Tonelli's Theorem implies $\lambda_\delta(Y_s) = 0$. \square

9. BINARY CODINGS OF REGIONS OF CO-ORIENTED HYPERPLANE ARRANGEMENTS AND THE GRADIENT VECTOR FIELD OF A RELU NEURAL NETWORK MAP

In this section, we collect elementary facts about co-oriented hyperplane arrangements that will be useful in the proofs of Theorems 5 and 6. We also introduce a partial orientation on $\mathcal{C}(F)_1$, the 1-skeleton of the bent hyperplane arrangement of a generic, transversal ReLU neural network, defined using the gradient of the neural network function F . This partially-oriented graph plays an important role in the proof of Theorem 6.

9.1. Regions, vertices, and edges of classical hyperplane arrangements.

Definition 9.1. *A region of a (possibly ordered, co-oriented) hyperplane arrangement \mathcal{A} in \mathbb{R}^n is a connected component of $\mathbb{R}^n \setminus \bigcup_{H \in \mathcal{A}} H$. Let $r(\mathcal{A})$ denote the number of regions of \mathcal{A} .*

Note that each region, R , of an ordered, co-oriented hyperplane arrangement $\mathbf{A} = \{\mathbf{H}_1, \dots, \mathbf{H}_k\}$ is naturally labeled with a binary k -tuple, $\vec{\theta} \in \{0, 1\}^k$, where the i^{th} component of $\vec{\theta}$ associated to R is 1 (resp. 0) if the co-orientation of \mathbf{H}_i points towards (resp., away from) R . We shall denote the region of \mathbf{A} labeled by the binary k -tuple $\vec{\theta}$ by $R_{\vec{\theta}}(\mathbf{A})$ and refer to it as the *$\vec{\theta}$ region of \mathbf{A}* . If the ordered, co-oriented hyperplane arrangement \mathbf{A} is clear from context, we will abbreviate the notation to $R_{\vec{\theta}}$. We will use $\bar{R}_{\vec{\theta}}(\mathbf{A})$ to denote the closure of $R_{\vec{\theta}}(\mathbf{A})$.

The assignment of binary k -tuples to regions of an ordered, co-oriented hyperplane arrangement \mathbf{A} is clearly injective, but it need not be surjective (see Figure 1). Lemma 9.2 gives a sufficient condition for the assignment to be bijective.

The contents of Lemmas 9.2 and 9.3 are well known and follow immediately from a classical theorem of Zaslavsky [27] (cf. [23, Thm. 2.5]).

Lemma 9.2. *Let $\mathbf{A} := \{\mathbf{H}_1, \dots, \mathbf{H}_k\}$ be a generic ordered, co-oriented hyperplane arrangement in \mathbb{R}^n , where $k \leq n$. Then $r(\mathbf{A}) = 2^k$, and hence there is a one-to-one correspondence between regions of \mathbf{A} and binary k -tuples.*

Lemma 9.3. *Let $\mathbf{A} = \{\mathbf{H}_1, \dots, \mathbf{H}_k\}$ be a (generic or non-generic) ordered, co-oriented hyperplane arrangement of $k \leq n$ hyperplanes in \mathbb{R}^n . \mathbf{A} has no bounded regions.*

Lemma 9.4. *Let $\mathcal{A} = \{H_1, \dots, H_N\}$ be any arrangement of N hyperplanes in \mathbb{R}^n . Then for any rank n , size n subarrangement $\{H_{i_1}, \dots, H_{i_n}\}$ of \mathcal{A} , $H_{i_1} \cap \dots \cap H_{i_n}$ is a 0-cell (vertex) of the canonical polyhedral complex $\mathcal{C}(\mathcal{A})$. Conversely, every 0-cell of $\mathcal{C}(\mathcal{A})$ can be realized as $H_{i_1} \cap \dots \cap H_{i_n}$ for some rank n , size n subarrangement $\{H_{i_1}, \dots, H_{i_n}\} \subseteq \mathcal{A}$.*

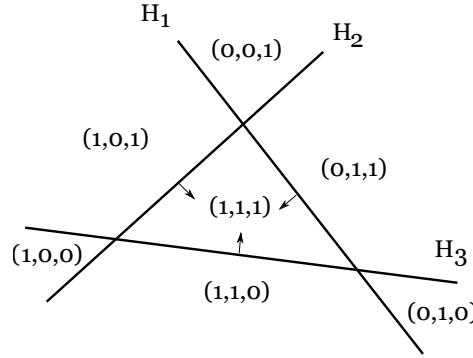


FIGURE 1. There are 7 regions in the complement of 3 co-oriented hyperplanes in \mathbb{R}^2 . The assignment of binary 3-tuples to regions is not surjective.

Proof. Let $\mathcal{A}' = \{H_{i_1}, \dots, H_{i_n}\}$ be a rank n , size n subarrangement of \mathcal{A} . From Lemma 2.5 it follows that \mathcal{A}' is a generic arrangement, and hence the n -fold intersection, $H_{i_1} \cap \dots \cap H_{i_n}$, is an affine subspace of \mathbb{R}^n of dimension 0. Indeed, since the $(n-k)$ -dimensional affine subspaces of \mathbb{R}^n associated to k -fold intersections of k -element subsets of \mathcal{A}' are reverse-ordered by inclusion, we see that $p = H_{i_1} \cap \dots \cap H_{i_n}$ is the unique 0-cell in the boundary of all cells of the polyhedral complex $\mathcal{C}(\mathcal{A}')$. Since \mathcal{A} is obtained from \mathcal{A}' by adding hyperplanes, $\mathcal{C}(\mathcal{A})$ is a polyhedral subdivision of $\mathcal{C}(\mathcal{A}')$, and so p is also a 0-cell (vertex) of $\mathcal{C}(\mathcal{A})$, as desired.

For the converse statement, we proceed by induction on n . For the base case ($n = 1$), it follows directly from the definition of $\mathcal{C}(\mathcal{A})$ that every 0-cell (vertex) is a hyperplane of \mathcal{A} . Now let $n > 1$ and suppose p is a 0-cell of $\mathcal{C}(\mathcal{A})$. We know that $p \in K$ for some hyperplane $K \in \mathcal{A}$. Consider the restricted *solution set arrangement*, $\mathcal{S}_K = \{K \cap H \mid H \in \mathcal{A} \setminus K\}$, from which we obtain a restricted hyperplane arrangement, \mathcal{A}_K , by deleting the degenerate solution sets. Then p is also a 0-cell in the canonical polyhedral complex $\mathcal{C}(\mathcal{A}_K)$. Since \mathcal{A}_K is an $(n-1)$ -dimensional hyperplane arrangement, the inductive hypothesis tells us that there exist hyperplanes $H_{i_1}, \dots, H_{i_{n-1}} \in \mathcal{A}$ such that $p = (K \cap H_{i_1}) \cap \dots \cap (K \cap H_{i_{n-1}})$ in K . Letting $H_{i_n} = K$, it follows that $p = H_{i_1} \cap \dots \cap H_{i_n}$. Moreover, \mathcal{A}' must be rank n , since otherwise its intersection would be an affine space of dimension > 0 . \square

Corollary 9.5. *Let $\mathcal{A} = \{H_1, \dots, H_N\} \subset \mathbb{R}^n, \{(W_i|b_i)\}_{i=1}^N$ be as above. There is a canonical surjective map from the set of linearly-independent n -element subsets (subbases), $\{W_{i_1}, \dots, W_{i_n}\} \subseteq \{W_1, \dots, W_N\}$, to the set of vertices of $\mathcal{C}(\mathcal{A})$.*

Proof. Subbases of the set $\{W_1, \dots, W_N\}$ of weight vectors are in canonical bijective correspondence with the set of rank n , size n subarrangements of \mathcal{A} . By Lemma 9.4, there is a surjective map from the set of rank n , size n subarrangements of \mathcal{A} to the set of vertices of $\mathcal{C}(\mathcal{A})$, defined by taking the n -fold intersection of the hyperplanes in the subarrangement. \square

Lemma 9.6. *Let $\mathcal{A} = \{H_1, \dots, H_N\} \subset \mathbb{R}^n, \{(W_i|b_i)\}_{i=1}^N$ be as above, and let*

$$\mathcal{B}_p = \{W_{i_1}, \dots, W_{i_n}\} \text{ and } \mathcal{B}_q = \{W_{j_1}, \dots, W_{j_n}\}$$

be two subbases of $\{W_1, \dots, W_N\}$ with corresponding canonical vertices p and q , respectively, as guaranteed by Corollary 9.5. If $|\mathcal{B}_p \cap \mathcal{B}_q| = n - 1$, then either $p = q$ or $p \cup q$ is the boundary of a 1-cell in $\mathcal{C}(\mathcal{A})$.

Proof. If

$$\mathcal{B}_p = \{W_{i_1}, \dots, W_{i_n}\} \text{ and } \mathcal{B}_q = \{W_{j_1}, \dots, W_{j_n}\}$$

are as above, then (reordering j_1, \dots, j_n if necessary) we may assume that $i_k = j_k$ for $k = 1, \dots, n - 1$ and $i_n \neq j_n$. Now (the proof of) Corollary 9.5 tells us that

$$p = (H_{i_1} \cap \dots \cap H_{i_{n-1}}) \cap H_{i_n} \text{ and } q = (H_{i_1} \cap \dots \cap H_{i_{n-1}}) \cap H_{j_n},$$

which tells us that p and q are points on the same 1-dimensional affine space, $(H_{i_1} \cap \dots \cap H_{i_{n-1}}) \subseteq \mathcal{C}(\mathcal{A})$. \square

Corollary 9.7. *Let $\mathcal{A} = \{H_1, \dots, H_{n+1}\}$ be any (generic or non-generic) arrangement of hyperplanes in \mathbb{R}^n . Every pair of 0-cells (vertices), $p \neq q$, of $\mathcal{C}(\mathcal{A})$ is connected by some 1-cell (edge) of $\mathcal{C}(\mathcal{A})$. That is, all vertices are adjacent in the graph $\mathcal{C}(\mathcal{A})_1$.*

Proof. Every pair of n -element subsets of an $(n+1)$ -element set has a common $(n-1)$ -element subset. So if $\mathcal{B}_p, \mathcal{B}_q$ are two subbases of $\{W_1, \dots, W_{n+1}\}$ and $p \neq q$, then Lemma 9.6 tells us p and q are adjacent in $\mathcal{C}(\mathcal{A})_1$. \square

The following lemma is immediate.

Lemma 9.8. *Let $\mathbf{A} = \{\mathbf{H}_1, \dots, \mathbf{H}_n\}$ be a generic, ordered, co-oriented arrangement of n hyperplanes in \mathbb{R}^n . Then the map from $\{0, 1\}^n$ to the set of faces of the polyhedral set $\bar{R}_{\vec{1}}$ given by*

$$(3) \quad \vec{\theta} \in \{0, 1\}^n \mapsto F_{\vec{\theta}} := \{\vec{\theta} \odot \vec{v} \mid \vec{v} \in \bar{R}_{\vec{1}}\}$$

is a bijection.

Here, $\vec{\theta} \odot \vec{v}$ denotes the *Hadamard product* (component-wise product) of $\vec{\theta}$ and \vec{v} . Accordingly, $F_{\vec{1}} = \bar{R}_{\vec{1}}$, and $F_{\vec{0}} = \{(0, \dots, 0)\}$. Moreover, $\dim(F_{\vec{\theta}}) = \sum_{i=1}^n \theta_i$.

Let \mathbf{A}^{st} be the standard ordered, co-oriented coordinate hyperplane arrangement in \mathbb{R}^n . That is, $\mathbf{A}^{st} = \{\mathbf{H}_i^{st}\}_{i=1}^n$, where

$$\mathbf{H}_i^{st} := \{\vec{v} = (v_1, \dots, v_n) \in \mathbb{R}^n \mid v_i = 0\},$$

co-oriented in the direction of the non-negative half-space, $\{\vec{v} \in \mathbb{R}^n \mid v_i \geq 0\}$. For $\vec{\theta} \in \{0, 1\}^n$, we shall denote by $F_{\vec{\theta}}^{st}$ the corresponding face of the polyhedral set $\bar{R}_{\vec{1}}^{st}$, as in (3). In summary, we have:

Lemma 9.9. *Let $\mathbf{A} = \{\mathbf{H}_1, \dots, \mathbf{H}_n\}$ be a generic, ordered, co-oriented arrangement of n hyperplanes in \mathbb{R}^n associated to a generic layer map $\sigma \circ A : \mathbb{R}^n \rightarrow \mathbb{R}^n$. Then A maps the \mathbf{A} decomposition of \mathbb{R}^n to the \mathbf{A}^{st} decomposition of \mathbb{R}^n , in the following sense:*

- $A(\mathbf{H}_i) = \mathbf{H}_i^{st}$ for all $i = 1, \dots, n$, and
- $A(R_{\vec{\theta}}) = R_{\vec{\theta}}^{st}$ and $A(F_{\vec{\theta}}) = F_{\vec{\theta}}^{st}$ for all $\vec{\theta} \in \{0, 1\}^n$.

Moreover, $\sigma \circ A$ is the composition of the affine isomorphism A realizing the above identification, followed by the projection $\bar{R}_{\vec{\theta}}^{st} \rightarrow F_{\vec{\theta}}^{st}$ given by taking the Hadamard product with $\vec{\theta} \in \{0, 1\}^n \subseteq \mathbb{R}^n$.

9.2. Regions of bent hyperplane arrangements and the ∇F -oriented 1-skeleton of the canonical polyhedral complex. We can similarly endow the activation regions of a generic, transversal ReLU neural network with a binary labeling as follows. Recall that if $F : \mathbb{R}^{n_0} \rightarrow \mathbb{R}$ is transversal and generic, Theorem 2 guarantees that the domain of the $(n_0 - 1)$ -skeleton of $\mathcal{C}(F)$ (resp., the n_0 -cells) agrees with the bent hyperplane arrangement, \mathcal{B}_F , (resp., the closures of the activation regions of \mathcal{B}_F).

In this case, the image of every activation region of F (interior of an n_0 -cell of $\mathcal{C}(F)$) is contained in a unique region of the co-oriented hyperplane arrangement in each layer:

Definition 9.10. *Let*

$$F : \mathbb{R}^{n_0} \xrightarrow{F_1} \mathbb{R}^{n_1} \xrightarrow{F_2} \dots \xrightarrow{F_m} \mathbb{R}^{n_m} \xrightarrow{G} \mathbb{R}$$

be a transversal, generic ReLU neural network, and let $\mathbf{A}^{(i)}$ denote the co-oriented hyperplane arrangement associated to F_i . The $(\vec{\theta}_1, \dots, \vec{\theta}_m)$ -region of F , denoted $R_{(\vec{\theta}_1, \dots, \vec{\theta}_m)}$, is the unique activation region of F satisfying the property that for each $i \in \{1, \dots, m\}$,

$$F_{i-1} \circ \dots \circ F_1 \left(R_{(\vec{\theta}_1, \dots, \vec{\theta}_m)} \right) \subseteq R_{\vec{\theta}_i} \left(\mathbf{A}^{(i)} \right).$$

Lemma 9.11. *Let*

$$F : \mathbb{R}^{n_0} \xrightarrow{F_1} \mathbb{R}^{n_1} \xrightarrow{F_2} \mathbb{R}^{n_2} \xrightarrow{F_3} \dots \xrightarrow{F_m} \mathbb{R}^{n_m} \xrightarrow{G} \mathbb{R}^1$$

be a transversal, generic ReLU neural network with associated weight matrices W_1, \dots, W_m, W_{m+1} . Let $R_{(\vec{\theta}_1, \dots, \vec{\theta}_m)}$ be a region of \mathcal{B}_F with associated sequence of binary tuples $(\vec{\theta}_1, \dots, \vec{\theta}_m)$, and let $\mathbf{p} \in \bar{R}$, where \bar{R} is the closure of R . Then

$$\left. \frac{\partial F}{\partial \mathbf{x}} \right|_{\mathbf{x}=\mathbf{p}} = W_{m+1} W_m^{\vec{\theta}_m} \cdots W_1^{\vec{\theta}_1},$$

where we define $W_k^{\vec{\theta}_k}$ to be the matrix obtained from W_k by replacing the i th row of W_k with 0's when the i th entry of $\vec{\theta}_k$ is 0.

Proof. Immediate from the definition of the affine-linear function on $R_{(\vec{\theta}_1, \dots, \vec{\theta}_m)}$ and the chain rule for partial derivatives. \square

The 1-skeleton, $\mathcal{C}(F)_1$, of the canonical polyhedral decomposition for a neural network $F : \mathbb{R}^{n_0} \rightarrow \mathbb{R}$, is an embedded linear graph in \mathbb{R}^{n_0} . This graph has a natural partial orientation, defined as follows.

Definition 9.12. *Let $F : \mathbb{R}^{n_0} \rightarrow \mathbb{R}$ be a neural network, denote by $\mathcal{C}(F)_1$ the 1-skeleton of its canonical polyhedral complex, and let C be a 1-cell in $\mathcal{C}(F)_1$.*

- *If F is nonconstant on C , orient C in the direction in which F increases.*
- *If F is constant on C , we will leave it unlabeled and refer to C as a flat edge.*

We will refer to this (partial) orientation on $\mathcal{C}(F)_1$ as the $\text{grad}(F)$ -orientation or ∇F -orientation.

Note that F is affine-linear on C by Theorem 1.

10. OBSTRUCTIONS TO TOPOLOGICAL EXPRESSIVENESS AND APPLICATIONS

This section uses the framework developed in the previous sections to give an alternative perspective on the Beise-Cruz-Schröder, Hanin-Sellke, Johnson result that a width n ReLU neural network $F : \mathbb{R}^n \rightarrow \mathbb{R}$ has decision regions that are either empty or unbounded. We also develop an architecture-based obstruction to the existence of multiple bounded connected components in a decision region.

Recall the statement of the Beise-Cruz-Schröder, Hanin-Sellke, Johnson result:

Theorem 5. [2, 15, 17] *Let $n \geq 2$. Suppose*

$$F : \mathbb{R}^n \xrightarrow{F_1} \mathbb{R}^n \xrightarrow{F_2} \dots \xrightarrow{F_m} \mathbb{R}^n \xrightarrow{G} \mathbb{R}$$

is a width n ReLU neural network map and $t \in \mathbb{R}$. Each of the sets $N_F(t), B_F(t), Y_F(t)$ (defined in (1)) is either empty or unbounded.

We will need some elementary facts about the image of a ReLU neural network in the width n case. We address the case of generic and non-generic layer maps separately.

Proposition 10.1. *Let*

$$N = (F_m \circ \dots \circ F_1) : \mathbb{R}^n \rightarrow \dots \rightarrow \mathbb{R}^n$$

be the composition of all but the final layer map of a generic width n ReLU neural network in which every hidden layer has dimension n . Then $\text{Im}(N) \subseteq \bar{R}_1^{st} \subseteq \mathbb{R}^n$ is the domain of a polyhedral complex, \mathcal{C} , with at most one cell of dimension n . Explicitly,

$$\text{Im}(N) = |\mathcal{C}| = \mathcal{P} \cup \mathcal{Q},$$

where \mathcal{P} is a (possibly empty) polyhedral set of dimension n , and \mathcal{Q} is a union of polyhedral sets of dimension $< n$. Moreover, if \mathcal{P} is nonempty, $n \geq 2$, and $\mathcal{P} = H_1^+ \cap \dots \cap H_m^+$ is an irredundant realization of \mathcal{P} as an intersection of closed half spaces, then $\{H_1, \dots, H_m\}$ has rank ≥ 2 .

Proof. We proceed by induction on the number, m , of layers. If $m = 1$, $\text{Im}(N)$ is $\bar{R}_{\vec{1}}^{st}$, which is the domain of a polyhedral complex with a single n -cell, $\mathcal{P} = \bar{R}_{\vec{1}}^{st}$, realizable as an irredundant intersection of n half-spaces. Moreover, if $n \geq 2$, the rank of the corresponding hyperplane arrangement, $\{H_1^{st}, \dots, H_n^{st}\}$ is ≥ 2 , as required.

Now consider $m \geq 2$ and suppose \mathcal{C}' is the polyhedral complex whose domain agrees with the image, $\text{Im}(N') = \mathcal{P}' \cup \mathcal{Q}'$, of the first $m - 1$ layers of N as described. Let

- \mathbf{A} be the ordered, co-oriented hyperplane arrangement associated to F_m ,
- $\mathcal{C}(\mathbf{A})$ be the associated polyhedral decomposition of the domain of F_m ,
- and $\mathcal{C}'_{\cap \mathcal{C}(\mathbf{A})}$ be the complex whose cells are pair-wise intersections of cells of \mathcal{C}' with cells of $\mathcal{C}(\mathbf{A})$.

By Lemma 9.9, F_m is affine-linear on the cells of $\mathcal{C}'_{\cap \mathcal{C}(\mathbf{A})}$. Indeed, we can alternatively characterize $\mathcal{C}'_{\cap \mathcal{C}(\mathbf{A})}$ as the level set complex, $\mathcal{C}'_{\in \mathcal{C}(\mathbf{A}^{st})}$, of F_m relative to the polyhedral decomposition associated to the standard hyperplane arrangement, \mathbf{A}^{st} , in the codomain.

Noting that the image of a polyhedral set under an affine-linear map is a polyhedral set, we now define \mathcal{C} to be the complex whose cells are the images of cells of $\mathcal{C}'_{\in \mathcal{C}(\mathbf{A}^{st})}$ under F_m .

It then follows immediately from Lemma 9.9 that all cells of \mathcal{C} have dimension $< n$ except possibly $F_m(\mathcal{P}' \cap \bar{R}_{\vec{1}}(\mathbf{A}))$. Further, since F_m is an affine isomorphism on $\bar{R}_{\vec{1}}(\mathbf{A})$, this cell will be n -dimensional iff $\mathcal{P}' \cap \bar{R}_{\vec{1}}(\mathbf{A})$ is n -dimensional.

In this case, we claim that as long as $n \geq 2$, an irredundant bounding hyperplane arrangement of the cell $\mathcal{P}' \cap \bar{R}_{\vec{1}}(\mathbf{A})$ will have rank ≥ 2 , since both \mathcal{P}' and $\bar{R}_{\vec{1}}(\mathbf{A})$ have this property. To see this, note that the union of the bounding hyperplane arrangements for \mathcal{P}' and $\bar{R}_{\vec{1}}(\mathbf{A})$ yields a bounding hyperplane arrangement for $\mathcal{P}' \cap \bar{R}_{\vec{1}}(\mathbf{A})$, and it necessarily has rank ≥ 2 . If this union is irredundant, we are done. If not, we appeal to Farkas' Lemma III, cf. [29, Sec. 2], which says that any *redundant* inequality in a system of linear inequalities is a non-negative linear combination of the other linear inequalities in the system. This implies that the rank of any bounding hyperplane arrangement of a polyhedral set is equal to the rank of an *irredundant* bounding hyperplane arrangement.

Moreover, the image under F_m of an irredundant bounding hyperplane arrangement for $\mathcal{P}' \cap \bar{R}_{\vec{1}}(\mathbf{A})$ is also irredundant and has rank ≥ 2 , since F_m is an affine isomorphism on $\bar{R}_{\vec{1}}(\mathbf{A})$. Defining \mathcal{P} to be

- $F_m(\mathcal{P}'_{\in \bar{R}_{\vec{1}}(\mathbf{A}^{st})})$ if the polyhedral set $\mathcal{P}'_{\in \bar{R}_{\vec{1}}(\mathbf{A}^{st})}$ has dimension n , and
- \emptyset otherwise,

and \mathcal{Q} to be the union of all other cells of \mathcal{C} , the result follows. \square

Lemma 10.2. *Let $N = (F_m \circ \dots \circ F_1) : \mathbb{R}^n \rightarrow \mathbb{R}^n$ be the composition of all but the final layer map of a generic width n ReLU neural network, and*

let $\text{Im}(N) = |\mathcal{C}| = \mathcal{P} \cup \mathcal{Q}$ as in Proposition 10.1 above. If $x \in |\mathcal{C}_{n-1}|$, then $N^{-1}(\{x\})$ is unbounded.

Proof. We proceed by induction on the number, m , of layers. The result is clear when $m = 1$, since in this case Lemma 9.9 tells us that $\mathcal{Q} = \emptyset$ and $\mathcal{P} = \bar{R}_{\vec{1}}^{st}$, and each point $x \in |\mathcal{C}_{n-1}| = \partial\mathcal{P} = \partial(\bar{R}_{\vec{1}}^{st})$ is in the image of the projection map $R_{\vec{\theta}}^{st} \rightarrow F_{\vec{\theta}}^{st}$, hence has unbounded preimage.

Now suppose $m > 1$ and the result holds for the image, $|\mathcal{C}'| = \mathcal{P}' \cup \mathcal{Q}'$, of the first $m-1$ layers of N . That is, each point in $|\mathcal{C}'_{n-1}|$ has unbounded preimage. Let $x \in |\mathcal{C}_{n-1}| \subseteq \text{Im}(N)$. If x is in the image of $|\mathcal{C}'_{n-1}|$, then it has unbounded preimage by the inductive hypothesis. So we may assume that x is in the image of $\text{int}(\mathcal{P}')$ and not in the image of $\partial\mathcal{P}'$. But since $x \in |\mathcal{C}_{n-1}|$, Lemma 9.9 implies that $F_m^{-1}(\{x\}) \subseteq |\mathcal{C}'|$ is a ray contained in $\text{int}(\mathcal{P}')$. Since $\text{int}(\mathcal{P}')$ is the image of the interior of a polyhedral set in the domain under a composition of affine-linear isomorphisms, the preimage of this ray is a ray in the domain, hence unbounded. The conclusion follows. \square

Lemma 10.3. *Let $\mathcal{P} = H_1^+ \cap \dots \cap H_m^+ \subseteq \mathbb{R}^n$ be an irredundant representation of a non-empty n -dimensional polyhedral set such that the hyperplane arrangement $\mathcal{A} = \{H_1, \dots, H_m\}$ has rank ≥ 2 , and let $F_i = \mathcal{P} \cap H_i$ be the bounding facet in $\partial\mathcal{P}$ corresponding to H_i . If $X \subseteq \mathbb{R}^n$ is any affine hyperplane satisfying $X \cap \mathcal{P} \neq \emptyset$, then $X \cap F_i \neq \emptyset$ for some i .*

Proof. Let X and \mathcal{P} in \mathbb{R}^n be as described above. Note that the rank assumption implies that $m, n \geq 2$. Since $X \cap \mathcal{P} \neq \emptyset$, then $X \subseteq \mathcal{P}$ or $X \cap \partial\mathcal{P} \neq \emptyset$. But if $X \subseteq \mathcal{P}$, then X must be parallel to every H_i , contradicting the rank assumption. So $X \cap \partial\mathcal{P} \neq \emptyset$. But $\partial\mathcal{P} = F_1 \cup \dots \cup F_m$, so $X \cap F_i \neq \emptyset$ for some i , as desired. \square

Lemma 10.4. *For any $n \in \mathbb{N}$, let $N = (F_m \circ \dots \circ F_1) : \mathbb{R}^n \rightarrow \mathbb{R}^n$ be the composition of all but the final layer map of a non-generic width n ReLU neural network. Then each point in $\text{Im}(N)$ has unbounded preimage.*

Proof. Let F_i be the first non-generic layer map in N . Since F_i is non-generic, the affine map A_i underlying F_i is non-invertible, by Lemma 2.5. Indeed, the preimage of any point is an affine-linear subspace of \mathbb{R}^n of dimension ≥ 1 , hence unbounded. Recalling that $F_i = \sigma \circ A_i$ is a map that is affine-linear on the cells of the canonical polyhedral decomposition $\mathcal{C}(F_{i-1} \circ \dots \circ F_i)$ of \mathbb{R}^n , it follows immediately that the affine-linear map on each cell is also non-invertible, hence the preimage of any point $p \in \text{Im}(F_i)$ is unbounded. Any point $q \in \text{Im}(N)$ is of the form $q = (F_m \circ \dots \circ F_{i+1})(p)$ for some $p \in \text{Im}(F_i)$. So $N^{-1}(\{q\})$ is unbounded. \square

Recall that a threshold t is transversal for a neural network F (Definition 6.6) if it is transversal for F with respect to its canonical polyhedral complex $\mathcal{C}(F)$.

Lemma 10.5. *Let t be a transversal threshold for a neural network $F : \mathbb{R}^{n_0} \rightarrow \mathbb{R}$. Then*

$$B_F(t) = \partial N_F(t) = \partial Y_F(t).$$

Proof. In the case that $B_F(t) = \emptyset$, either $Y_F(t) = \mathbb{R}^{n_0}$ and $N_F(t) = \emptyset$ or $Y_F(t) = \emptyset$ and $N_F(t) = \mathbb{R}^{n_0}$, so the statement holds.

Now suppose $B_F(t) \neq \emptyset$. For each cell $C \in \mathcal{C}(F)$ such that $F^{-1}(\{t\}) \cap C \neq \emptyset$, Lemma 5.4 guarantees that $\text{aff}(F^{-1}(\{t\}) \cap C)$ is a hyperplane that cuts C . Denote by C^+ and C^- the intersections of C with the two open half-spaces that are the complement of this hyperplane. Since F is nonconstant on C (by Lemma 5.3), precisely one of C^+, C^- must be contained in $Y_F(t)$ and the other must be contained in $N_F(t)$. Therefore $B_F(t) \subseteq \partial Y_F(t)$ and $B_F(t) \subseteq \partial N_F(t)$. The reverse inclusions are obvious. \square

We are now ready for:

Proof of Theorem 5. Fix an integer $n \geq 2$ and let $F : \mathbb{R}^n \rightarrow \mathbb{R}$ be a ReLU neural network whose hidden layers all have dimension $\leq n$. We may then assume WLOG that every intermediate layer has dimension n (padding points in layers of dimension $< n$ with 0s as needed). Decompose F as $F = G \circ N$, where $N : \mathbb{R}^n \rightarrow \mathbb{R}^n$ is the composition of all the layer maps except for the final one, and $G : \mathbb{R}^n \rightarrow \mathbb{R}$ is the final layer map. If G is degenerate, it is immediate that the theorem holds. So assume that G is nondegenerate.

Step 1: We will prove that for every $t \in \mathbb{R}$, the decision boundary $B_F(t)$ is either empty or unbounded. Since G is nondegenerate, $X_t := G^{-1}(\{t\})$ is an affine hyperplane in the final hidden layer of F . Note that

$$B_F(t) = N^{-1}(G^{-1}(\{t\})) = N^{-1}(\text{Im}(N) \cap X_t).$$

If $\text{Im}(N) \cap X_t$ is empty, then $B_F(t)$ is empty, as desired. So assume $\text{Im}(N) \cap X_t$ is nonempty.

Case 1: We first consider the case that N is non-generic. In this case, $\text{Im}(N) \cap X_t \neq \emptyset$ implies $N^{-1}(\text{Im}(N) \cap X_t)$ is unbounded by Lemma 10.4.

Case 2: Now consider the case that N is generic. By Proposition 10.1, $\text{Im}(N)$ is the domain of a polyhedral complex \mathcal{C} that has a unique (possibly empty) n -cell \mathcal{P} .

Subcase a: If $X_t \cap \mathcal{P} = \emptyset$, then the assumption that $\text{Im}(N) \cap X_t \neq \emptyset$ implies $X_t \cap \mathcal{C}_{n-1} \neq \emptyset$. Therefore $B_F(t)$ is unbounded, by Lemma 10.2.

Subcase b: If $X_t \cap \mathcal{P} \neq \emptyset$, then \mathcal{P} is nonempty. Because $n \geq 2$, Proposition 10.1 guarantees that \mathcal{P} has rank ≥ 2 . Hence $X_t \cap \mathcal{C}_{n-1} \neq \emptyset$ by Lemma 10.3. Therefore $B_F(t)$ is unbounded by Lemma 10.2.

Step 2: We will use the fact that $B_F(t)$ is empty or unbounded to show $N_F(t)$ and $Y_F(t)$ are also.

Case 1: When $t \in \mathbb{R}$ is a transversal threshold, it is now straightforward to see that the decision regions $Y_F(t)$ and $N_F(t)$ are also either empty or unbounded, since $B_F(t) = \partial N_F(t) = \partial Y_F(t)$ by Lemma 10.5, and a bounded set cannot have unbounded closure.

Case 2: Suppose $t \in \mathbb{R}$ is a non-transversal threshold. We will give an argument for $Y_F(t)$; the argument for $N_F(t)$ is analogous. Let X_t^+ be the positive half-space associated to the co-oriented affine hyperplane X_t . Then

$$F^{-1}((t, \infty)) = N^{-1}(\text{Im}(N) \cap X_t^+).$$

If this intersection is empty, then $Y_F(t)$ is empty, as desired. If this intersection is nonempty, there are two subcases.

Subcase a: $t < \max\{t' \in \mathbb{R} : F^{-1}((t', \infty)) \neq \emptyset\}$. In this case, since F is continuous, there exists $\epsilon > 0$ such that

$$(t, t + \epsilon) \subset \{t' \in \mathbb{R} : F^{-1}((t', \infty)) \neq \emptyset\}.$$

Hence Lemma 5.5 implies the existence of a transversal threshold $t' > t$ for which $F^{-1}(\{t'\})$ is non-empty. Noting that $F^{-1}(\{t'\}) = B_F(t')$, it follows from the first part of the proof that $B_F(t')$ is unbounded. So $Y_F(t) \supseteq B_F(t')$ must also be unbounded.

Subcase b: $t = \max\{t' \in \mathbb{R} : F^{-1}((t', \infty)) \neq \emptyset\}$. In this case $Y_F(t)$ is empty, as desired. \square

10.1. Obstructing multiple bounded connected components. As observed in [17], it is straightforward to construct, for every n , a width $n + 1$ neural network with a single hidden layer, $\mathbb{R}^n \rightarrow \mathbb{R}^{n+1} \rightarrow \mathbb{R}$, that has a bounded decision region consisting of a single connected component. We prove that such a simple architecture cannot produce a decision region with more than one bounded connected component.

Theorem 6. *Let $F : \mathbb{R}^n \rightarrow \mathbb{R}^{n+1} \rightarrow \mathbb{R}$ be a ReLU neural network. Then a decision region $Y_F(t)$ or $N_F(t)$ associated to a transversal threshold t can have no more than one bounded connected component.*

The following Lemma is a standard result in linear programming.

Lemma 10.6. *Let $\mathcal{P} \subseteq \mathbb{R}^n$ be a polyhedral set, and let $F : \mathcal{P} \rightarrow \mathbb{R}$ be an affine-linear map on \mathcal{P} . If F achieves a maximum (resp., minimum) on the interior of \mathcal{P} , then F achieves this maximum (resp., minimum) value on all of \mathcal{P} , and hence on all faces in its boundary.*

Corollary 10.7. *Let $F : \mathbb{R}^{n_0} \rightarrow \mathbb{R}$ be a neural network, and $\mathcal{C}(F)$ be its canonical polyhedral complex. If \mathcal{P} is a cell of $\mathcal{C}(F)$ with at least one vertex as a face, and F achieves a maximum (resp., minimum) on \mathcal{P} , then F achieves a maximum (resp., minimum) at a vertex of \mathcal{P} .*

Proof. The function F is affine-linear on \mathcal{P} by Theorem 1. Under the given assumptions, every face of \mathcal{P} will also be a polyhedral set with a vertex as a face. The result follows by strong induction on the dimension of \mathcal{P} , applying Lemma 10.6 in the inductive step. \square

Corollary 10.8. *Let F and $\mathcal{C}(F)$ be as above. If \mathcal{P} is a bounded cell (polytope) of $\mathcal{C}(F)$ of any dimension, then F achieves a maximum (resp., minimum) at a vertex of \mathcal{P} .*

Proof. Every polytope $\mathcal{P} \subset \mathbb{R}^n$ (of any dimension) has at least one vertex. Moreover, it is bounded, hence compact, since cells are closed. The extreme value theorem then guarantees that F achieves both a minimum and maximum value on \mathcal{P} . The result follows from Corollary 10.7. \square

Proposition 10.9. *Let t be a transversal threshold for a neural network $F : \mathbb{R}^{n_0} \rightarrow \mathbb{R}$, and let S be a bounded connected component of $Y_F(t)$ (resp., $N_F(t)$). Then there exist non-empty bounded subgraphs $\mathcal{G}' \subseteq \mathcal{G} \subseteq \mathcal{C}(F)_1$ which, when endowed with the ∇F -orientation (Definition 9.12), satisfy:*

- (i) \mathcal{G}' is flat;

- (ii) $\mathcal{G}' \subseteq \mathcal{G} \subsetneq S$;
- (iii) there is a non-empty collection, \mathcal{E} , of edges adjacent to \mathcal{G} , satisfying the property that every edge $e \in \mathcal{E}$ points towards \mathcal{G} (resp., points away from \mathcal{G}) and has nonempty intersection with ∂S and the other decision region, $N_F(t)$ (resp., $Y_F(t)$).

Remark 10.10. One can view the graph \mathcal{G}' described in Proposition 10.9 as a piecewise linear analogue of a Morse critical point of index n (resp., 0).

Proof. Let \overline{S} denote the closure of S . Since \overline{S} is closed and bounded, hence compact, the extreme value theorem tells us that F attains its maximum (resp., minimum) value, $M \in \mathbb{R}$ (resp., $m \in \mathbb{R}$), on \overline{S} . I.e., there exists $x \in \overline{S}$ such that $F(x) = M$ and $F(y) \leq M$ for all $y \in \overline{S}$. But Lemma 10.6 implies that $F^{-1}(\{M\})$ contains a non-empty subgraph, \mathcal{G}' , of $\mathcal{C}(F)_1$, since a maximum value, if attained on the interior of a cell, is attained on the whole cell, including its boundary.

Moreover, $\mathcal{G}' \subset S$, for if $\mathcal{G}' \cap (\partial S \subseteq F^{-1}\{t\}) \neq \emptyset$ then $t = M$, which would imply that t is *not* a transversal threshold since its preimage contains a vertex which by definition cannot have a nonconstant cellular neighborhood.

Let \mathcal{G} be the maximal subgraph of $\mathcal{C}(F)_1 \cap S$ that both contains \mathcal{G}' and is entirely contained in S . Properties (i) and (ii) are immediate by construction. To see Property (iii), note that $\mathcal{C}(F)_1$ is connected and unbounded, so it follows that $\mathcal{C}(F)_1 \cap \partial S \neq \emptyset$. Since t is a transversal threshold, all points in $\partial \overline{S}$ have nonconstant cellular neighborhood, hence all edges of $\mathcal{C}(F)_1$ intersecting ∂S are oriented, and the orientations are toward (resp., away from) \mathcal{G} if $S \subseteq Y_F(t)$ (resp., $S \subseteq N_F(t)$). \square

Proof of Theorem 6. Since F has a single hidden layer, $\mathcal{C}(F) = \mathcal{C}(\mathcal{A})$, for the hyperplane arrangement, $\mathcal{A} \subseteq \mathbb{R}^n$ associated to the first layer map. We may assume without loss of generality that $|\mathcal{A}| = n + 1$, for if the first layer map is degenerate then the neural network has width n , and hence its decision regions have no bounded connected components, by Theorem 5.

Now let $t \in \mathbb{R}$ be a transversal threshold for a (not necessarily generic or transversal) ReLU network, $F : \mathbb{R}^n \rightarrow \mathbb{R}^{n+1} \rightarrow \mathbb{R}$. We will show that $Y_F(t)$ has no more than one bounded connected component. The argument for $N_F(t)$ is analogous.

Assume, aiming for a contradiction, that $Y_F(t)$ has more than one bounded connected component. Choose two of these: S_1 and S_2 . As described in Proposition 10.9, there exist non-empty bounded subgraphs $\mathcal{G}_i \subset S_i$ of the 1-skeleton of $\mathcal{C}(\mathcal{A}) = \mathcal{C}(F)$ and associated non-empty collections, \mathcal{E}_i , of edges adjacent to \mathcal{G}_i , equipped with ∇F -orientation pointing towards \mathcal{G}_i . For each S_i , choose an *external* vertex $p_i \in \mathcal{G}_i$. That is, choose a vertex $p_i \in \mathcal{G}_i$ in the boundary of an edge of \mathcal{E}_i .

By Corollary 9.7, p_1 and p_2 are connected by a 1-cell $e \in \mathcal{C}(\mathcal{A})$. It follows that e is in both \mathcal{E}_1 and \mathcal{E}_2 . But this is impossible, since it would require e to be oriented in two different directions at once. We conclude that one of \mathcal{E}_1 or \mathcal{E}_2 must be empty, hence Proposition 10.9 tells us that one of S_1, S_2 must be empty. The result follows. \square

BOSTON COLLEGE; DEPARTMENT OF MATHEMATICS; 522 MALONEY HALL; CHESTNUT HILL, MA 02467

Email address: grigsbyj@bc.edu

BOSTON COLLEGE; DEPARTMENT OF MATHEMATICS; 567 MALONEY HALL; CHESTNUT HILL, MA 02467

Email address: lindseka@bc.edu