

PIERCE–BIRKHOFF CONJECTURE IS TRUE FOR SPLINES

ZEHUA LAI AND LEK-HENG LIM

ABSTRACT. We prove the Pierce–Birkhoff conjecture for splines, i.e., continuous piecewise polynomials of degree d in n variables on a hyperplane partition of \mathbb{R}^n , can be written as a finite lattice combination of polynomials. We will provide a purely existential proof, followed by a more in-depth analysis that yields effective bounds.

1. INTRODUCTION

The Pierce–Birkhoff conjecture states that every continuous piecewise polynomial function on a semialgebraic partition of \mathbb{R}^n is a finite lattice combination of polynomials, i.e., the functions generated by addition, multiplication, and maximization of polynomial functions on \mathbb{R}^n . To be precise, a continuous function $f : \mathbb{R}^n \rightarrow \mathbb{R}$ that takes a form

$$(1) \quad f = g_1 \mathbb{1}_{\Pi_1} + \cdots + g_r \mathbb{1}_{\Pi_r}$$

where $g_1, \dots, g_r \in \mathbb{R}[x_1, \dots, x_n]$,

$$(2) \quad \mathbb{R}^n = \Pi_1 \cup \cdots \cup \Pi_r$$

a partition, and $\mathbb{1}_{\Pi}$ the indicator function, is a *spline* of degree $d := \max\{\deg g_1, \dots, \deg g_r\}$ if (2) is a hyperplane partition, i.e., Π_1, \dots, Π_r are polytopes. This is the near universal definition of splines in both pure [3, 4, 29] and applied mathematics [8, 9, 33]. It is equivalent to requiring that (2) be a triangulation, i.e., Π_1, \dots, Π_r are simplices, since any hyperplane partition may be refined into a triangulation and vice versa. It also includes all smoother C^k -variants as special cases, i.e., the results that we prove in this article for will hold for all C^k -splines, $k = 0, 1, 2, \dots$

There is however a generalization called *semialgebraic splines* [12, 13, 30] that allows for Π_1, \dots, Π_r to be semialgebraic sets. In this terminology, the Pierce–Birkhoff conjecture states:

Conjecture 1.1 (Pierce–Birkhoff). *Any semialgebraic spline $f : \mathbb{R}^n \rightarrow \mathbb{R}$ can be expressed as*

$$(3) \quad f = \max_{i=1, \dots, p} \min_{j=1, \dots, p'} f_{ij}$$

with finitely many polynomials $f_{ij} \in \mathbb{R}[x_1, \dots, x_n]$, $i = 1, \dots, p$, $j = 1, \dots, p'$.

Despite its apparent simplicity and longevity (first in-print appearance in 1962), the conjecture has resisted all attempts so far. Aside from the case $d = 1$ and the case $n \leq 2$, the Pierce–Birkhoff conjecture is not known to hold for any other values of d and n . We will prove that it holds for splines, i.e., when the partition (2) is a hyperplane partition, for all d and n .

Brief history and prior attempts. The Pierce–Birkhoff conjecture, as noted in [23], was first formally formulated by Henriksen and Isbell in [15], with roots in Birkhoff and Pierce’s earlier work [5] in 1956. It has been shown to be true for $n = 1$ and 2 in [24] and for $d = 1$ in [28]. These results have spawned many variations: For $n = 1$ it has been extended to nonsingular curves [22] and semialgebraic curves [26]. For $n = 2$ it has been extended to arbitrary ordered fields [10], nonsingular two-dimensional real affine varieties [32], regular local rings of dimension two [19]; it has also been restricted to the hypercube [18] and generalized to signomials [11]. Nevertheless we emphasize that none of these variations go beyond $n = 2$. There are only limited partial results for

the $n = 3$ case [25, 21]; the latter work [21] is a partial result for the $n = 3$ case of the connectedness conjecture, which is known to imply the Pierce–Birkhoff conjecture [20].

Our contributions. We will show that Conjecture 1.1 holds for any hyperplane partition. We first provide a short existential proof that uses the language of real spectrum [6] and the notion of separating ideals [22]. We then give a more careful construction to deduce explicit bounds for p, p' and the degrees of f_{ij} 's in terms of r, d , and n .

2. SPLINES AND SEMIALGEBRAIC SPLINES

We begin with a more formal description of partitions and splines. For any $b \in \mathbb{N}$, we define the set of all 3^b ternary bits by

$$(4) \quad \Theta_b := \{\theta : \{1, \dots, b\} \rightarrow \{1, 0, -1\}\}.$$

Definition 2.1 (Partitions). Let $P := \{\pi_1, \dots, \pi_b\} \subseteq \mathbb{R}[x_1, \dots, x_n]$ be a set of nonconstant polynomials of degrees not more than d' . For any $\theta \in \Theta_b$, we define

$$\Pi_\theta := \{x \in \mathbb{R}^n : \text{sgn}(\pi_i(x)) = \theta(i), i = 1, \dots, b\}.$$

Then $\{\Pi_\theta : \theta \in \Theta_b\}$ is called a *semialgebraic partition* of \mathbb{R}^n of degree d' induced by π_1, \dots, π_b [6]. If $d' = 1$, i.e., π_1, \dots, π_b are all degree-one polynomials, then it is called a *hyperplane partition* [1]. We will call each Π_θ a *piece* of the partition. In a semialgebraic partition, any piece Π_θ has a well-defined dimension [6].

By far most common and classical partition is a triangulation, i.e., a subdivision of \mathbb{R}^n into a finite simplicial complex [14] or one may also use convex polytopes in place of simplices [8, 9, 33]. The important point to note is that any such “piecewise linear” partitions may be refined into a hyperplane partition as in Definition 2.1, even if they are not at first defined in this manner.

A partition in the sense of Definition 2.1 uniquely determines a partition in the sense of (2) with $r = 3^b$, i.e.,

$$\mathbb{R}^n = \bigcup_{\theta \in \Theta_b} \Pi_\theta = \Pi_1 \cup \dots \cup \Pi_r.$$

Whenever there is no need to concern ourselves with the defining polynomials π_1, \dots, π_b , we will label a partition simply as $\mathbb{R}^n = \Pi_1 \cup \dots \cup \Pi_r$.

Definition 2.2 (Chambers). Let $\{\Pi_\theta : \theta \in \Theta_b\}$ be a semialgebraic partition of \mathbb{R}^n . A piece of full dimension n is called a *chamber* [1] if $\{\Pi_\theta : \theta \in \Theta_b\}$ is a hyperplane partition and more generally a *semialgebraic chamber*.

We may now formally define splines and semialgebraic splines.

Definition 2.3 (Splines). Let $\{\Pi_\theta : \theta \in \Theta_b\}$ be semialgebraic partition of \mathbb{R}^n of degree d' . A continuous function $f : \mathbb{R}^n \rightarrow \mathbb{R}$ is a *semialgebraic spline* of degree d if f restricts to a polynomial $g_\theta \in \mathbb{R}[x_1, \dots, x_n]$ of degree not more than d on Π_θ whenever $\Pi_\theta \neq \emptyset$. A semialgebraic spline on a hyperplane partition, i.e., $d' = 1$, is called a *spline*.

To be clear, we have written

$$d := \max\{\deg g_\theta : \theta \in \Theta_b\}, \quad d' := \max\{\deg \pi_i : i = 1, \dots, b\}$$

to distinguish the two different degrees tied to a semialgebraic spline.

We will be invoking this simple observation several times later: Given a semialgebraic partition $\mathbb{R}^n = \Pi_1 \cup \dots \cup \Pi_r$, assume without loss of generality that Π_1, \dots, Π_p , $p \leq r$, are all the chambers in this partition. Let V_i be the (Euclidean) closure of Π_i , $i = 1, \dots, p$. Then $\mathbb{R}^n = V_1 \cup \dots \cup V_p$. By continuity, the restriction of f to Π_i completely determines f on V_i , $i = 1, \dots, p$. A consequence is that g_{p+1}, \dots, g_r , i.e., restrictions of f to non-chamber pieces, do not play a role in our results.

Definition 2.4 (Max-definable function). A *max-definable functions* in the variables x_1, \dots, x_n is a function $f : \mathbb{R}^n \rightarrow \mathbb{R}$ generated by indeterminates x_1, \dots, x_n and constants $c \in \mathbb{R}$ under four operations: addition $(x, y) \mapsto x + y$, multiplication $(x, y) \mapsto x \cdot y$, maximization $(x, y) \mapsto \max(x, y)$, and scalar multiplication $(c, x) \mapsto cx$.

Note that minimization comes for free in Definition 2.4 as $\min(x, y) := -\max(-x, -y)$. The identity

$$x \max(y, 0) = \max(\min(xy, x^2y + y), \min(0, -x^2y - y))$$

allows us to exchange the order of multiplication and maximization using this rule. Consequently any max-definable functions can be written in the form $\max_{i=1, \dots, p} \min_{j=1, \dots, p'} g_{ij}$ as in the statement of Pierce–Birkhoff conjecture [15]. We will prove that any spline is max-definable, i.e., Conjecture 1.1 is true for any d and n when $d' = 1$.

3. EXISTENTIAL PROOF

The proof in this section is based on the notion of separating ideals [22]. We will present a simplified treatment defined for two semialgebraic sets U, V , which is all we need, after recalling some real algebraic geometry [6, 22].

The real spectrum of a ring R , denoted $\text{Spec}_{\mathbb{R}}(R)$, is a set of pairs $(\mathfrak{p}_\alpha, >_\alpha) =: \alpha$, where \mathfrak{p}_α is a prime ideal and $>_\alpha$ is a total order on R/\mathfrak{p}_α [6, Proposition 7.1.2]. For any $f \in R$, let

$$(5) \quad f(\alpha) := f + \mathfrak{p}_\alpha \in R/\mathfrak{p}_\alpha.$$

If $\mathfrak{p}_\alpha \subseteq \mathfrak{p}_\beta$ and the natural map $R/\mathfrak{p}_\alpha \rightarrow R/\mathfrak{p}_\beta$ preserves order, we write $\alpha \subseteq \beta$ and say that β *specializes* α . The topology on $\text{Spec}_{\mathbb{R}}(R)$ is generated by basic open sets of the form

$$\tilde{B}(f) = \{\alpha \in \text{Spec}_{\mathbb{R}}(R) : f(\alpha) >_\alpha 0\}$$

over all $f \in R$. Under this topology, β specializes α if and only if β is in the closure of $\{\alpha\}$ [6, Proposition 7.1.18].

Henceforth, we will limit our attention to $R := \mathbb{R}[x_1, \dots, x_n]$. We identify $a \in \mathbb{R}^n$ with its vanishing ideal \mathfrak{m}_a and let the unique order on \mathfrak{m}_a be given by $f < g$ if and only if $f(a) < g(a)$. This defines an injection $\mathbb{R}^n \rightarrow \text{Spec}_{\mathbb{R}}(R)$ and we may identify \mathbb{R}^n with its image in $\text{Spec}_{\mathbb{R}}(R)$. Given a semialgebraic set V defined by a Boolean combination of sets of the form

$$B(f) = \{a \in \mathbb{R}^n : f(a) > 0\}$$

over some collection of $f \in R$, the same Boolean combination of $\tilde{B}(f)$ defines a set \tilde{V} in $\text{Spec}_{\mathbb{R}}(R)$. We also set $\tilde{\emptyset} = \emptyset$. Clearly $\tilde{V} \cap \mathbb{R}^n = V$. In fact \tilde{V} is uniquely determined by V and does not depend on the choice of Boolean expressions used to define V [6, Proposition 7.2.2].

For any $f \in R$, the polynomial function $f : \mathbb{R}^n \rightarrow \mathbb{R}$ extends to a function $f : \text{Spec}_{\mathbb{R}}(R) \rightarrow \mathbb{R}$ as in (5). Note that $f \geq 0$ on V if and only if $f \geq 0$ on \tilde{V} . Hence a semialgebraic partition of \mathbb{R}^n gives a partition of $\text{Spec}_{\mathbb{R}}(R)$. A semialgebraic spline f may then be extended to $\text{Spec}_{\mathbb{R}}(R)$ in a natural way: If f is as in (1), then it can be extended to $\text{Spec}_{\mathbb{R}}(R)$ as

$$f = g_1 \mathbb{1}_{\tilde{\Pi}_1} + \dots + g_r \mathbb{1}_{\tilde{\Pi}_r},$$

where $g_i : \text{Spec}_{\mathbb{R}}(R) \rightarrow \mathbb{R}$ extends the polynomial function g_i , $i = 1, \dots, r$. We will write $f_\alpha \in R$ for a polynomial that agrees with f at $\alpha \in \text{Spec}_{\mathbb{R}}(R)$. The choice of f_α will not affect any result in this article. Lastly, by virtue of the way they are defined, addition, multiplication, and maximum of two semialgebraic splines over \mathbb{R}^n extend naturally over $\text{Spec}_{\mathbb{R}}(R)$.

We will define separating ideal as in [22] and introduce a slight generalization.

Definition 3.1 (Separating ideal). Let $\alpha, \beta \in \text{Spec}_{\mathbb{R}}(R)$. We write $\langle \alpha, \beta \rangle$ for the ideal generated by all $f \in R$ such that $f(\alpha) \geq 0$ and $f(\beta) \leq 0$. This is called the separating ideal of α and β . Let

U, V be semialgebraic sets. Likewise we will write $\langle U, V \rangle$ for the ideal of R generated by all $f \in R$ such that $f \geq 0$ on U and $f \leq 0$ on V , and call it the separating ideal of U and V .

Since $f \geq 0$ on U if and only if $f \geq 0$ on \tilde{U} for any $f \in R$, the latter definition extends unambiguously to $\langle \tilde{U}, \tilde{V} \rangle$. It is clear that if $\alpha \in \tilde{U}$ and $\beta \in \tilde{V}$, then $\langle U, V \rangle \subseteq \langle \alpha, \beta \rangle$. For easy reference, we reproduce the main result of [22].

Theorem 3.2 (Madden, 1989). *A semialgebraic spline f is max-definable if and only if for any $\alpha, \beta \in \text{Spec}_{\mathbb{R}}(R)$, $f_{\alpha} - f_{\beta} \in \langle \alpha, \beta \rangle$.*

This alternative formulation of Pierce–Birkhoff conjecture has come to be known as the local Pierce–Birkhoff conjecture [21]. While it sheds light on the original conjecture, it does not appear to be any easier. Neither the separating ideal $\langle \alpha, \beta \rangle$ nor functions of the form $f_{\alpha} - f_{\beta}$ are well understood. For instance, it is known that the separating ideal of disjoint semialgebraic sets $\langle U, V \rangle$ can be zero: Take $U = \{(x, y) \in \mathbb{R}^2 : xy \geq 1, x > 0\}$ and $V = \mathbb{R}^2 \setminus U$, then $\langle U, V \rangle = \{0\}$. There are also examples of dimension three [7]. But there is no known example of the form $\langle \alpha, \beta \rangle$, which would be a potential counterexample for the local Pierce–Birkhoff conjecture.

Nevertheless, Theorem 3.2 is useful to us in the form of the following corollary.

Corollary 3.3. *Let f be a semialgebraic spline over a semialgebraic partition $\mathbb{R}^n = \Pi_1 \cup \dots \cup \Pi_r$ with chambers Π_1, \dots, Π_p , $p \leq r$. Let $g_i \in \mathbb{R}[x_1, \dots, x_n]$ be the restriction of f to Π_i , and V_i be the closure of Π_i , $i = 1, \dots, p$. If $g_i - g_j \in \langle V_i, V_j \rangle$, then f is max-definable.*

Proof. Given any $\alpha, \beta \in \text{Spec}_{\mathbb{R}}(R)$, we need to verify the condition in Theorem 3.2. Since V_1, \dots, V_p cover \mathbb{R}^n , $\tilde{V}_1, \dots, \tilde{V}_p$ cover $\text{Spec}_{\mathbb{R}}(R)$. For any α, β , we can find $\alpha \in \tilde{V}_i, \beta \in \tilde{V}_j$, then $f_{\alpha} - f_{\beta} = g_i - g_j \in \langle V_i, V_j \rangle \subseteq \langle \alpha, \beta \rangle$. Hence f is max-definable. \square

While Corollary 3.3 has not been explicitly stated anywhere, the proof of Pierce–Birkhoff conjecture for $n = 2$ [24] essentially uses this simplified version. By a further refinement if necessary, we may assume that the partition $\mathbb{R}^n = \Pi_1 \cup \dots \cup \Pi_r$ gives a cylindrical algebraic decomposition [6]. For $n = 2$, the intersection $V_i \cap V_j$ is either empty, a point, or a curve. It is relatively straightforward to study how two sets in the plane can intersect in these three cases and arrive at the Pierce–Birkhoff conjecture. However, for $n \geq 3$, the only easy case is when $V_i \cap V_j$ is a hypersurface with the separating ideal a principal ideal. When codimension exceeds one, the separating ideal becomes far more complicated.

Fortunately, if we limit ourselves to splines as in Definition 2.3, then the intersections are all convex polytopes, which brings the condition in Corollary 3.3 within reach of standard convex geometry.

Theorem 3.4 (Pierce–Birkhoff is true for splines). *Let f be a spline. Then f is max-definable.*

Proof. Let f be a spline as in Definition 2.3. We follow the notations in Corollary 3.3. Let chambers in the partition be Π_1, \dots, Π_p , and V_i be the closure of Π_i , $i = 1, \dots, p$. For any two V_i, V_j . The intersection $V_i \cap V_j$ is a closed convex polytope. Let \mathfrak{a} be the vanishing ideal of the affine span of $V_i \cap V_j$. We claim that $\langle V_i, V_j \rangle = \mathfrak{a}$.

We first consider the case $V_i \cap V_j \neq \emptyset$. One direction is clear: If $h \in \langle V_i, V_j \rangle$, then $h = 0$ in the closed convex set $V_i \cap V_j$, so $h = 0$ in its affine span, thus $\langle V_i, V_j \rangle \subseteq \mathfrak{a}$.

Conversely, choose a point x in the relative interior of $V_i \cap V_j$. Let C_i and C_j be the tangent cones of V_i and V_j at x . Then $C_i \cap C_j$ is the affine span of $V_i \cap V_j$ and it is the largest affine subset that C_i or C_j contains. Consider the quotients $C_i/(C_i \cap C_j)$ and $C_j/(C_i \cap C_j)$, or, alternatively, the projections of C_i and C_j to the orthogonal complement of $C_i \cap C_j$. These are two closed disjoint pointed polyhedral cones and thus there is a hyperplane $\{x \in \mathbb{R}^n : h(x) = 0\}$ that strictly separates the two cones [16]. By definition, the degree-one polynomial $h \in \langle V_i, V_j \rangle$. For any degree-one polynomial $h' \in \mathfrak{a}$, the hyperplane $\{x \in \mathbb{R}^n : h'(x) = 0\}$ contains $C_i \cap C_j$. So $h + \varepsilon h'$ will still

separate the two cones for sufficiently small $\varepsilon > 0$. Therefore $h + \varepsilon h' \in \langle V_i, V_j \rangle$ and thus $h' \in \langle V_i, V_j \rangle$. Hence $\mathfrak{a} \subseteq \langle V_i, V_j \rangle$. In conclusion, $\mathfrak{a} = \langle V_i, V_j \rangle$.

If $V_i \cap V_j = \emptyset$, there is also a strictly separating hyperplane [31, Theorem 1.1.10]. Hence we can find h so that $h + \varepsilon$ separates V_i and V_j for all sufficiently small $\varepsilon > 0$. Thus $1 \in \langle V_i, V_j \rangle$, i.e., $\langle V_i, V_j \rangle = R$. In this case the affine span of $V_i \cap V_j$ is the empty set and we also get $\langle V_i, V_j \rangle = \mathfrak{a}$.

By its definition, a spline f has $f|_{V_i} = g_i$, $i = 1, \dots, p$. Since $g_i - g_j = 0$ is in $V_i \cap V_j$, $g_i - g_j \in \mathfrak{a} = \langle V_i, V_j \rangle$. By Corollary 3.3, f is max-definable. \square

We will see in Theorem 4.2 that Theorem 3.4 can be made much more informative, but at some cost to clarity. The abstract approach above, freed of the need to keep track of polynomials, is more likely to generalize to other semialgebraic splines.

4. EFFECTIVE BOUNDS

In this section, we will derive explicit bounds for Theorem 3.4, essentially by making effective the abstract constructions in [22].

Lemma 4.1. *Let f be a semialgebraic spline over a semialgebraic partition $\mathbb{R}^n = \Pi_1 \cup \dots \cup \Pi_r$ with chambers Π_1, \dots, Π_p , $p \leq r$. Let $g_i \in \mathbb{R}[x_1, \dots, x_n]$ be the restriction of f to Π_i , and V_i be the closure of Π_i , $i = 1, \dots, p$. If for $i, j = 1, \dots, p$, there exists $h_{ij} \in \mathbb{R}[x_1, \dots, x_n]$ such that $h_{ij} \leq g_j - g_i$ on V_j and $h_{ij} \geq 0$ on V_i , then f is max-definable with*

$$f = \max_{i=1, \dots, p} \min_{j=1, \dots, p} f_{ij}$$

for some $f_{ij} \in \mathbb{R}[x_1, \dots, x_n]$ where

$$\max\{\deg f_{ij} : i, j = 1, \dots, p\} \leq \max\{\deg g_i, \deg h_{ij} : i, j = 1, \dots, p\}.$$

Proof. For any $i, j = 1, \dots, p$, let

$$f_{ij} := h_{ij} + g_i.$$

So the assumption becomes $f_{ij} \leq g_j$ on V_j and $f_{ij} \geq g_i$ on V_i . For $i = j$, the two conditions on h_{ij} forces $h_{ii} = 0$, so $f_{ii} = g_i$. Set

$$f_0 := \max_{i=1, \dots, p} \min_{j=1, \dots, p} f_{ij}.$$

On V_k , for any $i \neq k$,

$$\min_{j=1, \dots, p} f_{ij} \leq f_{jk} \leq g_k,$$

and for $i = k$,

$$\min_{j=1, \dots, p} f_{ij} = f_{ii} = g_k.$$

So

$$f_0 = \max_{i=1, \dots, p} \min_{j=1, \dots, p} f_{ij} = g_k$$

on V_k . Since $f|_{V_k} = g_k$, $k = 1, \dots, p$, we have $f_0 = f$ and the degree bound follows. \square

We now establish our final result.

Theorem 4.2 (Effective Pierce–Birkhoff for splines). *Let f be a spline over a hyperplane partition $\mathbb{R}^n = \Pi_1 \cup \dots \cup \Pi_r$ induced by degree-one polynomials $P = \{\pi_1, \dots, \pi_b\}$. Let Π_1, \dots, Π_m , $m \leq r$, be the chambers, $g_i \in \mathbb{R}[x_1, \dots, x_n]$ be the restriction of f to Π_i , and V_i be the closure of Π_i , $i = 1, \dots, m$. If $d = \max\{\deg g_i : i = 1, \dots, m\}$, then f is max-definable with*

$$f = \max_{i=1, \dots, p} \min_{j=1, \dots, p} f_{ij}$$

for some $f_{ij} \in \mathbb{R}[x_1, \dots, x_n]$, $\deg f_{ij} \leq 2d + 1$, and

$$p = O(b^{2n^2} d^n)$$

for any fixed $n \in \mathbb{N}$.

Proof. Let $i, j \in \{1, \dots, m\}$. Let \mathfrak{a} be the vanishing ideal of the affine span of $V_i \cap V_j$. As we saw in the proof of Theorem 3.4, $\mathfrak{a} = \langle V_i, V_j \rangle$. When $V_i \cap V_j \neq \emptyset$, \mathfrak{a} is an ideal generated by degree-one polynomials and $g_i - g_j$ can then be written as

$$(6) \quad g_i - g_j = \sum_{k=1}^q \pi_{ijk} \rho_{ijk}$$

where $\deg \pi_{ijk} = 1$, $\deg \rho_{ijk} \leq \deg(g_i - g_j) - 1$, and $q := q_{ij}$ is understood to depend on i, j but denoted as q to avoid clutter (we will later see that $q_{ij} \leq n + 1$ irrespective of i, j). Furthermore, the construction in the proof of Theorem 3.4 shows that we may choose π_{ijk} so that they separate V_i and V_j . When $V_i \cap V_j = \emptyset$, $g_i - g_j$ can similarly be expressed as in (6), where π_{ijk} are degree-one separating polynomials and $\deg \rho_{ijk} \leq \deg(g_i - g_j)$. Now we add these polynomials to P to obtain

$$(7) \quad P' = P \cup \{g_i - g_j, \rho_{ijk}, \rho_{ijk} - 1 : i, j = 1, \dots, m, k = 1, \dots, q\}.$$

Then P' induces a finer semialgebraic (but not necessarily hyperplane) partition Π' . Let

$$\{V_{i\hat{i}} : i = 1, \dots, m, \hat{i} = 1, \dots, p_i\}$$

be an exhaustive list of the closures of chambers in Π' so that we have

$$\mathbb{R}^n = \bigcup_{i=1}^m V_i \quad \text{and} \quad V_i = \bigcup_{\hat{i}=1}^{p_i} V_{i\hat{i}}, \quad i = 1, \dots, m.$$

Let $p := p_1 + \dots + p_m$. This refined partition Π' has the property that any $g_i - g_j$ and any ρ_{ijk} and $\rho_{ijk} - 1$ is either nonnegative or nonpositive on any $V_{i\hat{i}}$.

Take any pair $V_{i\hat{i}}$ and $V_{j\hat{j}}$. Suppose $i = j$. Then $f|_{V_{i\hat{i}}} = f|_{V_{j\hat{j}}} = g_i$. So 0 satisfies the condition in Lemma 4.1 that $0 \leq g_i - g_i$ on $V_{i\hat{i}}$ and $0 \geq 0$ on $V_{j\hat{j}}$. Suppose $i \neq j$. If $g_i - g_j \geq 0$ on V_i , then 0 satisfies the condition that $0 \leq g_i - g_j$ on $V_{i\hat{i}}$ and $0 \geq 0$ on $V_{j\hat{j}}$. It remains to consider the case $g_i - g_j \leq 0$ on V_i . Write $g_i - g_j$ as in (6). We may choose the sign of π_{ijk} and ρ_{ijk} so that $\pi_{ijk} \geq 0$ on $V_{i\hat{i}}$, and $\pi_{ijk} \leq 0$ on $V_{j\hat{j}}$. Now define

$$\sigma_{ijk} := \begin{cases} 0 & \rho_{ijk} \leq 0 \text{ on } V_{i\hat{i}}; \\ \rho_{ijk} & \rho_{ijk} \geq 0 \text{ on } V_{i\hat{i}}, \rho_{ijk} \geq 0 \text{ on } V_{j\hat{j}}; \\ 1 & 0 \leq \rho_{ijk} \leq 1 \text{ on } V_{i\hat{i}}, \rho_{ijk} \leq 0 \text{ on } V_{j\hat{j}}; \\ \rho_{ijk}^2 & \rho_{ijk} \geq 1 \text{ on } V_{i\hat{i}}, \rho_{ijk} \leq 0 \text{ on } V_{j\hat{j}}. \end{cases}$$

It is straightforward to check that $g_i - g_j \leq \sum_{k=1}^q \pi_{ijk} \sigma_{ijk}$ on $V_{i\hat{i}}$, $\sum_{k=1}^q \pi_{ijk} \sigma_{ijk} \geq 0$ on $V_{j\hat{j}}$, and

$$\deg \left(\sum_{k=1}^q \pi_{ijk} \sigma_{ijk} \right) \leq 2 \deg(g_i - g_j) + 1 \leq 2d + 1.$$

Hence by Lemma 4.1, f can be written as

$$f = \max_{i=1, \dots, p} \min_{\hat{i}=1, \dots, p} f_{i\hat{i}}$$

where $f_{i\hat{i}} \in \mathbb{R}[x_1, \dots, x_n]$ has the desired degree bound $2d + 1$.

To obtain the bound for p , we need to count the number of chambers in the semialgebraic partition Π' . We started from the hyperplane partition Π induced by P . Since $|P| = b$, the partition Π is cut out by b hyperplanes and it can have at most $O(b^n)$ chambers [27, Chapter 6]. So $m = O(b^n)$. The refined semialgebraic partition Π' is induced by P' in (7). There are $m^2 = O(b^{2n})$ polynomials of the form $g_i - g_j$ in P' . For each fixed $i, j \in \{1, \dots, m\}$, we may choose π_{ijk} in (6) so that the number of summands $q = q_{ij} \leq n + 1$. To see this, first suppose $V_i \cap V_j \neq \emptyset$. Let $\dim(V_i \cap V_j) = n_{ij} \leq n$ and let $h_1, \dots, h_{n-n_{ij}}$ generate $\langle V_i, V_j \rangle$. Then we may choose π_{ijk} to be of the forms $h, h + \varepsilon h_1, \dots, h + \varepsilon h_{n-n_{ij}}$ where $\varepsilon > 0$ is a fixed constant as in the proof of Theorem 3.4. Thus $q_{ij} \leq n - n_{ij} + 1 \leq n + 1$. Now suppose $V_i \cap V_j = \emptyset$. Then we may choose π_{ijk} to be of the two

forms $h, h + \varepsilon$. Thus $q_{ij} \leq 2 \leq n + 1$. Hence there are at most $2m^2(n + 1) = O(b^{2n})$ polynomials of the form ρ_{ijk} or $\rho_{ijk} - 1$ in P' . All included, we obtain

$$|P'| \leq b + m^2 + 2m^2(n + 1) = O(b) + O(b^{2n}) + O(b^{2n}) = O(b^{2n}).$$

Since the polynomials in P' have degrees $d' \leq 2d$, it follows from [2, Chapter 7] that p , the number of semialgebraic chambers in Π' , is bounded by

$$O((|P'| \cdot d')^n) = O((b^{2n} \cdot 2d)^n) = O(b^{2n^2} d^n)$$

as required. \square

A slightly more careful count would yield an explicit expression for the leading constant c_n , which depends only on n , in the $O(b^{2n^2} d^n)$ bound. We see little point in providing it since but the bound for p is already exponential in n .

5. CONCLUSION

While our goal in this article is purely mathematical, there is one real-world consequence that is perhaps worth highlighting. By [17, Theorem 3.8], it follows from Theorem 3.4 that every spline is a ReLU-activated transformer.

REFERENCES

- [1] M. Aguiar and S. Mahajan. *Topics in hyperplane arrangements*, volume 226 of *Mathematical Surveys and Monographs*. American Mathematical Society, Providence, RI, 2017.
- [2] S. Basu, R. Pollack, and M.-F. Roy. *Algorithms in real algebraic geometry*, volume 10 of *Algorithms and Computation in Mathematics*. Springer-Verlag, Berlin, 2003.
- [3] L. J. Billera. Homology of smooth splines: generic triangulations and a conjecture of Strang. *Trans. Amer. Math. Soc.*, 310(1):325–340, 1988.
- [4] L. J. Billera. The algebra of continuous piecewise polynomials. *Adv. Math.*, 76(2):170–183, 1989.
- [5] G. Birkhoff and R. S. Pierce. Lattice-ordered rings. *An. Acad. Brasil. Ci.*, 28:41–69, 1956.
- [6] J. Bochnak, M. Coste, and M.-F. c. Roy. *Real algebraic geometry*, volume 36 of *Ergebnisse der Mathematik und ihrer Grenzgebiete (3)*. Springer-Verlag, Berlin, 1998.
- [7] L. Bröcker. On the separation of basic semialgebraic sets by polynomials. *Manuscripta Math.*, 60(4):497–508, 1988.
- [8] C. K. Chui. *Multivariate splines*, volume 54 of *CBMS-NSF Regional Conference Series in Applied Mathematics*. Society for Industrial and Applied Mathematics, Philadelphia, PA, 1988.
- [9] C. de Boor, K. Höllig, and S. Riemenschneider. *Box splines*, volume 98 of *Applied Mathematical Sciences*. Springer-Verlag, New York, 1993.
- [10] C. N. Delzell. On the Pierce–Birkhoff conjecture over ordered fields. volume 19, pages 651–668. 1989. Quadratic forms and real algebraic geometry (Corvallis, OR, 1986).
- [11] C. N. Delzell. Extension of the two-variable Pierce–Birkhoff conjecture to generalized polynomials. *Ann. Fac. Sci. Toulouse Math. (6)*, 19:37–56, 2010.
- [12] M. DiPasquale and F. Sottile. Bivariate semialgebraic splines. *J. Approx. Theory*, 254:105392, 19, 2020.
- [13] M. DiPasquale, F. Sottile, and L. Sun. Semialgebraic splines. *Comput. Aided Geom. Design*, 55:29–47, 2017.
- [14] A. Hatcher. *Algebraic Topology*. Cambridge University Press, 2002.
- [15] M. Henriksen and J. R. Isbell. Lattice-ordered rings and function rings. *Pacific J. Math.*, 12:533–565, 1962.
- [16] V. L. Klee, Jr. Separation properties of convex cones. *Proc. Amer. Math. Soc.*, 6:313–318, 1955.
- [17] Z. Lai, L.-H. Lim, and Y. Liu. Attention is a smoothed cubic spline. *arXiv:2408.09624*, 2024.
- [18] S. Lapenta and I. Leuştean. Towards understanding the Pierce–Birkhoff conjecture via MV-algebras. *Fuzzy Sets and Systems*, 276:114–130, 2015.
- [19] F. Lucas, J. Madden, D. Schaub, and M. Spivakovsky. Approximate roots of a valuation and the Pierce–Birkhoff conjecture. *Ann. Fac. Sci. Toulouse Math. (6)*, 21(2):259–342, 2012.
- [20] F. Lucas, J. J. Madden, D. Schaub, and M. Spivakovsky. On connectedness of sets in the real spectra of polynomial rings. *Manuscripta Math.*, 128(4):505–547, 2009.
- [21] F. Lucas, D. Schaub, and M. Spivakovsky. On the Pierce–Birkhoff conjecture. *Journal of Algebra*, 435:124–158, 2015.
- [22] J. J. Madden. Pierce–Birkhoff rings. *Arch. Math. (Basel)*, 53(6):565–570, 1989.
- [23] J. J. Madden. Henriksen and Isbell on f -rings. *Topology Appl.*, 158(14):1768–1773, 2011.

- [24] L. Mahé. On the Pierce–Birkhoff conjecture. *Rocky Mountain J. Math.*, 14(4):983–985, 1984.
- [25] L. Mahé. On the Pierce–Birkhoff conjecture in three variables. *J. Pure Appl. Algebra*, 211(2):459–470, 2007.
- [26] M. Marshall. The Pierce–Birkhoff conjecture for curves. *Canad. J. Math.*, 44(6):1262–1271, 1992.
- [27] J. Matoušek. *Lectures on discrete geometry*, volume 212 of *Graduate Texts in Mathematics*. Springer-Verlag, New York, 2002.
- [28] S. Ovchinnikov. Max-min representation of piecewise linear functions. *Beiträge Algebra Geom.*, 43(1):297–302, 2002.
- [29] D. Plaumann. Positivity of continuous piecewise polynomials. *Bull. Lond. Math. Soc.*, 44(4):749–757, 2012.
- [30] B. Shekhtman and T. Sorokina. A note on intrinsic supersmoothness of bivariate semialgebraic splines. *Comput. Aided Geom. Design*, 98:Paper No. 102137, 5, 2022.
- [31] J. Stoer and C. Witzgall. *Convexity and optimization in finite dimensions I*, volume 163. Springer Science & Business Media, 2012.
- [32] S. Wagner. On the Pierce–Birkhoff conjecture for smooth affine surfaces over real closed fields. *Ann. Fac. Sci. Toulouse Math. (6)*, 19:221–242, 2010.
- [33] R.-H. Wang. *Multivariate spline functions and their applications*, volume 529 of *Mathematics and its Applications*. Kluwer Academic Publishers, Dordrecht, 2001.

DEPARTMENT OF MATHEMATICS, UNIVERSITY OF TEXAS, AUSTIN, TX 78712
 Email address: zehua.lai@austin.utexas.edu

COMPUTATIONAL AND APPLIED MATHEMATICS INITIATIVE, DEPARTMENT OF STATISTICS, UNIVERSITY OF CHICAGO, CHICAGO, IL 60637
 Email address: lekheng@uchicago.edu