

Master Thesis in Mathematics

Spring Semester 2025

Philipp Nazari

The Geometry of Generalization

Supervisors: Clemens Hutter, Prof. Dr. Helmut Bölskei

October 2025

Abstract

While recent developments undoubtedly demonstrate the power of deep learning, we still lack a fundamental understanding of why overparameterized models work so well in practice. A common explanation attributes this phenomenon to implicit regularization induced by first-order optimization techniques like SGD. However, recent work has found that even zeroth-order guess-and-check optimizers very frequently find well generalizing minima. In this work, we mathematically formulate this heuristic, known as the volume hypothesis. We then fully establish existing research ideas which, using a tropical geometric perspective, introduce a dual representation of fully connected feedforward ReLU networks. This abstraction offers a perspective for studying the volume hypothesis which, to the best of our knowledge, is novel. While deriving general results remains challenging, we analyze multiple lower-dimensional examples, some inspired by Telgarsky’s sawtooth construction, which support the volume hypothesis. In particular, using the tropical geometric framework, we argue that exponentially complex minima in the loss landscape are unstable, leading learning algorithms to converge to solutions where the network does not fully utilize its available expressivity. Our work provides a novel perspective to think about generalization of deep ReLU networks, and we hope to inspire further theoretical and empirical research to establish more general results. The code for this project can be found at <https://github.com/phnazari/geomgen>.

Acknowledgments

I would like to express my gratitude to Prof. Bölcskei for offering me the chance to explore such an interesting area of deep learning theory. I'm also deeply thankful to Clemens Hutter for his thoughtful supervision. Your guidance and knowledgeable input enabled me to dive headfirst into tropical geometry and deep ReLU networks.

Finally, I don't even know where to start thanking my parents. You have always supported me on my journey through life, providing me with the perfect support to explore my academic curiosity.

Contents

Abstract	i
Acknowledgments	ii
Notation	v
1 Introduction	1
1.1 Perspectives on ReLU Networks	4
1.2 Complexity of ReLU Networks	4
2 Mathematical Background	6
2.1 Neural Networks	6
2.1.1 Fully Connected Feedforward Networks	6
2.1.2 Classification vs. Regression	8
2.1.3 Training Neural Networks	9
2.2 On Sets, Functions and (In)Dependence	9
2.3 Polyhedral Complexes	14
3 Affine Geometry	18
3.1 Affine and (D)CPA Functions	18
3.2 Affine Dualities	19
3.3 CPA Functions as Upper Convex Hulls	21
3.4 Tessellations	23
4 Tropical Geometry	26
4.1 Basic Definitions	26
4.2 Relation to Affine Geometry	28
5 Dual Representation of Neural Networks	30
5.1 Neural Networks and Affine Geometry	30
5.1.1 Positive and Negative Samples	34
5.1.2 Counting Dual Points	35
5.2 Neural Networks and Tropical Geometry	39
6 Characterizing the Decision Boundary	46
6.1 Identifying Cells and Faces	46

6.2	Application to the Decision Boundary	50
7	Characterizing Affine Regions	55
7.1	Counting Affine Regions	60
8	Telgarsky’s Sawtooth Network	66
8.1	Problem Setting	66
8.2	The Dual Representation	68
8.3	Experiments	74
9	One Layer – Random Weights	78
9.1	Summing Circles	79
9.2	Summing Rotated Circles	83
9.3	Putting It All Together	87
9.4	Implications	102
10	Heuristic Arguments in Dual Space	114
10.1	Large Upper Convex Hulls	114
10.2	Gaussian Dual Points	118
11	Conclusion	123
A	Helping Lemmas	128

Notation

The following table contains a (non-exhaustive) selection of the most frequently used notation. We provide links to formal definitions for non generic symbols.

Notation	Description
\mathbb{N}	The set of natural numbers $\{1, 2, 3, \dots\}$
\mathbb{N}_0	The set of natural numbers including 0
$[m : n]$	The set $\{m, m + 1, \dots, n\}$ for $m, n \in \mathbb{N}_0$
\mathbf{v}	Multi-dimensional vector
\mathbf{v}^T	Transpose of a vector \mathbf{v}
(\mathbf{x}, y)	Point in \mathbb{R}^{d+1} with $\mathbf{x} \in \mathbb{R}^d$ and $y \in \mathbb{R}$
$\langle \cdot, \cdot \rangle$	Euclidean inner product
$\ \cdot \ $	Norm of a vector or function
$\ \cdot \ _2$	Euclidean norm
$\xrightarrow{\sim}$	Bijection
$ \cdot $	Cardinality of a set
\mathbf{A}	Matrix
\mathbf{A}_i	i 'th row of matrix \mathbf{A}
x^+	$\max(0, x)$
x^-	$\max(0, -x)$
$\mathcal{N}(\mu, \sigma^2)$	Gaussian distribution
$U(S)$	Uniform distribution over the set S
H_n	n 'th Harmonic number $\sum_{i=1}^n \frac{1}{i}$
$f \sim g$	The functions f and g are asymptotically equivalent, $\lim_{n \rightarrow \infty} \frac{f(n)}{g(n)} = 1$
$g = \mathcal{O}(f)$	There exist constants $C > 0$ and $N \in \mathbb{N}$ such that $ g(n) \leq C f(n) $ for all $n \geq N$
$g = \Theta(n)$	There exist constants $c_1, c_2 > 0$ and $N \in \mathbb{N}$ such that $c_1 f(n) \leq g(n) \leq c_2 f(n)$ for all $n \geq N$
\sqcup	Disjoint union
\boxplus	Sum of a scalar and a set of vectors, Definition 5.1.5
\oplus	Tropical addition, Definition 4.1.1
\odot	Tropical multiplication, Definition 4.1.1
\oslash	Tropical quotient, Definition 4.1.6
\boxplus	Minkowski sum
$f_{\mathbf{a},b}$	Affine map $f_{\mathbf{a},b}(\mathbf{x}) = \langle \mathbf{a}, \mathbf{x} \rangle + b$
$\rho_t(\cdot)$	The function $\max(\cdot, t)$
\mathcal{N}	Fully connected ReLU network, see Definition 2.1.1

c	Complexity measure for neural network \mathcal{N} , for example number of affine regions or linear pieces in decision boundary
\mathcal{B}	Decision boundary of a binary classification network, Definition 2.1.7
L	Depth of neural network
d	Input dimension for neural network
n_i	Width of layer $i = 1, \dots, L$ of neural network
∇	Gradient operator
$(\mathbf{x}, y) \in f$	(\mathbf{x}, y) lies on the graph of f , Definition 2.2.1
$(\mathbf{x}, y) \succ f$	(\mathbf{x}, y) lies above the graph of f , Definition 2.2.1
$\text{affhul}(X)$	Affine hull of X , the smallest affine space containing X
$\mathcal{C}(X)$	Convex hull of X
$\mathcal{U}(X)$	Upper convex hull of X , Definition 2.3.4
$\mathcal{U}_k(X)$	k -skeleton of $U(X)$, Definition 2.3.4
$\mathcal{U}^*(X)$	Upper convex hull vertices of X , Definition 2.3.4
$f \mathcal{U}(X)$	Affine function f is tangent to the upper convex hull of X , Definition 6.1.5
CPA	Convex and piecewise affine function
DCPA	Difference of convex and piecewise affine functions
CPA(d)	Set of CPA functions $\mathbb{R}^d \rightarrow \mathbb{R}$
DCPA(d)	Set of DCPA functions $\mathbb{R}^d \rightarrow \mathbb{R}$
$\mathcal{Q}(S)$	CPA function induced by a set S of dual points, Definition 3.3.1
$\mathcal{T}(F)$	Tessellation induced by CPA function F , Definition 3.4.1
$\mathcal{T}_k(F)$	k -skeleton of $\mathcal{T}(F)$, Definition 3.4.3
$\mathcal{T}(S)$	Tessellation induced by a CPA function $\mathcal{Q}(S)$
$\mathcal{T}(P, N)$	Tessellation induced by the DCPA function $\mathcal{Q}(P) - \mathcal{Q}(N)$, Definition 3.4.5
(P_l, N_l)	Dual representation of $\mathcal{Q}(P_L) - \mathcal{Q}(N_L)$ up to layer l , Corollary 5.1.8
σ	Cell in $\mathcal{T}(F)$
ζ	Face in $\mathcal{Q}(S)$
$\mathfrak{P}(P \diamond N)$	Set of paths of dual points, Definition 7.1.6
$\mathfrak{P}(P, N)$	Set of paths of d -cells, Definition 7.1.4
$ \Sigma $	Support of a polyhedral complex Σ , Definition 2.3.5
\mathfrak{R}	Real space, Page 19
\mathfrak{D}	Dual space, Page 19
$\text{Aff}_{\mathfrak{R}}(d)$	Real affine space, Page 19
$\text{Aff}_{\mathfrak{D}}(d)$	Dual affine space, Page 19
\mathcal{R}	Bijection between \mathfrak{D} and $\text{Aff}_{\mathfrak{R}}(d)$, Lemma 3.2.1
$\check{\mathcal{R}}$	Bijection between \mathfrak{R} and $\text{Aff}_{\mathfrak{D}}(d)$, Lemma 3.2.3
S_I	Given an index-set I and an indexed set S , $S_I := \{s_i \mid i \in I\}$
$\mathbf{A}_=$	Set of implicit equality constraints of a polynomial $\{\mathbf{Ax} \geq \mathbf{b}\}$, Definition 2.3.7
\mathbf{A}^σ	Inequality constraints defining a cell σ , Remark 6.1.2
f_m	Mirror-map, Definition 8.1.3

Chapter 1

Introduction

Conventional learning theory predicts overparameterized deep neural networks to overfit the training data. Yet, that effect is not observed in practice [1, 2]. This phenomenon was long attributed to implicit regularization induced by first order optimization techniques like SGD, which bias the network towards well-generalizing minima [3, 4, 5]. However, Chiang et al. [6] show, using a zeroth order optimization technique, that this implicit regularization is not necessary for finding such minima. Instead, they argue that the generalization ability depends solely on the structure of the loss-landscape.

Valle-Pérez et al. [7] heuristically link generalization to the simplicity of a network, arguing that our physical universe is “simple”, and therefore any training algorithm should favor simple hypotheses over complex ones in order to accurately capture the true rule of the universe (i.e., to generalize). In addition to this heuristic argument, it seems intuitive that a large capacity allows deep models to overfit to noise, leading to poor generalization error. The tendency of learning algorithms to prefer simple optima, even when the model could overfit the data, is called “simplicity bias” [2, 8, 9, 10].

More specifically, the maximum number of affine regions of a ReLU networks is known to grow exponentially in depth and polynomially in width [11, 12, 13, 14]. Telgarsky [15] further demonstrated the representational benefits of depth by constructing a deep, narrow network that achieves an exponential (in the number of layers) number of affine regions. This construction enables correct classification of a hand-crafted dataset that any shallow network that is not very wide fails to classify.

This raises the question: do all networks exhibit an exponential number of affine regions? To answer this question, it is important to note that Telgarsky achieved exponential expressivity through a carefully designed recursive architecture. In practice, Hanin and Rolnick [16] show that, at initialization, the number of affine regions grows linearly with the number of neurons along any one-dimensional subspace, and thus polynomially per volume of input-space [14]. Crucially, they demonstrate that neural networks, in practice, do not realize the full exponential complexity theoretically available (further discussed in Section 1.2).

In this work we propose a new perspective on the simplicity bias, and thus on generalization capabilities of deep networks, using a geometric argument. Specifically, we study the “Volume Hypothesis” [6], which states that well generalizing minima occupy larger volumes (i.e., are more flat) in parameter space than badly generalizing ones.

Before diving into more details on our contribution, we introduce the volume hypothesis more rigorously. By randomly sampling networks until finding one with 100% training accuracy, Chiang et al. [6] show (in the setting of binary classification) that the majority of the resulting networks are simple and generalize well. This leads them to attribute the simplicity bias solely to the geometry of the loss landscape, independent of the first order regularization induced by SGD. They coin their observation the **Volume Hypothesis**:

“[...] generalizing minima occupy a much larger volume than poorly generalizing minima in neural loss functions, and [...] this volume disparity alone is enough to explain generalization [...]” [6, p.2].

In order to make this argument mathematically more rigorous, assume the network parameters are drawn i.i.d. from a probability distribution, in practice a Gaussian. Given a dataset $D \subseteq \mathbb{R}^d$, let E be the event that all training samples are classified correctly by a network \mathcal{N} .

Let furthermore \mathfrak{c} be a complexity measure for \mathcal{N} (e.g., the number of linearly pieces in the decision boundary of a ReLU classifier or the number of affine regions in a ReLU regression network).

To study the volume hypothesis, we are interested in the posterior density $\mathbb{P}(\mathfrak{c}|E = 1)$ of the complexity given that the network achieves 100% training accuracy. Using Bayes’ rule, this density can be re-written as

$$\begin{aligned} \mathbb{P}(\mathfrak{c}|E = 1) &= \frac{\mathbb{P}(E = 1|\mathfrak{c})\mathbb{P}(\mathfrak{c})}{\mathbb{P}(E)} \\ &\sim \mathbb{P}(\mathfrak{c})\chi_{E=1}. \end{aligned}$$

In other words, the posterior density of the complexity at a minimum of the loss landscape is proportional to the volume of the complexity in parameter space.

In this work, we focus our attention on fully connected, feedforward ReLU networks, which allows leveraging tropical geometry (Section 2.3.1) to derive a *dual representation* of the network. This representation enables us to identify the networks with the upper convex hull of two sets of points (“dual points”, Proposition 5.1.7). As a result, instead of directly analyzing deep ReLU networks, we can shift our focus to better understanding upper convex hulls and their complexity (Chapters 8-10).

As a complexity measure, we use the number of affine regions in the setting of regression (as in [11, 12, 13, 17]) and the number of linear pieces in the decision boundary in the setting of classification (as in [18]). Those two quantities can be directly tied to the dual representation: the number of linear pieces in the decision boundary corresponds to the number of a specific kind of edge in an upper convex hull (Theorem 6.1.7), and

the number of linear pieces corresponds to the number of vertices in another upper convex hull (Theorem 7.0.5).

Abstracting away from ReLU networks to upper convex hulls using the dual representation, the distribution of weights and biases induces a distribution of dual points. However, the distribution of the dual points is generally intractable, as the marginal distributions of each point are highly interdependent. To address this issue, we focus on studying a variety of simpler, low-dimensional problems. For instance, in Chapter 8 we derive the dual representation of Telgarsky’s construction [15]. In Chapter 9, we investigate the effect of adding a one-dimensional ReLU layer with Gaussian weights and biases to an existing deterministic network, using dual representations inspired by our previous analysis of Telgarsky’s sawtooth network. Here, we show that the probability of increasing the networks complexity by adding the random layer decreases as the starting network becomes more complex (Proposition 9.4.3). Additionally, we show that the expected marginal gain in complexity is expected to decrease (Corollary 9.4.13).

The primary contribution of this work is to provide strong evidence for the volume hypothesis in simpler, low-dimensional toy settings. To achieve this, we establish tropical geometry as a new setting to understanding generalization. Additionally, we rigorously derive duality results, filling gaps in existing proofs and arguments.

Overall, the structure of our work is as follows. In Chapter 2, we begin by providing fundamental knowledge and relevant definitions, covering topics from standard neural network theory to polyhedral complexes, probability and statistics. Chapter 3 introduces affine geometry, which is then linked to tropical geometry in Chapter 4. In Chapter 5, we establish a duality result between ReLU networks and upper convex hulls, which is then applied in Chapter 6 to characterize the decision boundary of binary classification networks, and in Chapter 7 to analyze affine regions. In Chapters 8-10, we apply the previously developed knowledge in a number simpler settings, providing more evidence for the volume hypothesis.

In summary, our contributions are as follows:

1. we connect the volume hypothesis to tropical geometry, effectively transferring the problem of investigating ReLU networks to the investigation of upper convex hulls of points,
2. we fully formalize the relationship between the complexity of ReLU networks and properties of their dual representations (Chapter 6 and Chapter 7),
3. we explicitly compute the dual representation of Telgarsky’s [15] sawtooth construction (Chapter 8),
4. we provide evidence for the volume-hypothesis in simplified settings (e.g., Chapter 9), demonstrating in a specific case that a network we conjecture to have exponential complexity (i.e., one that lives in the exponential complexity regime, see Section 1.2) is expected to transition to the subexponential complexity regime,

5. we study additional toy examples throughout Chapter 10.

1.1 Perspectives on ReLU Networks

It is well known [11, 13, 14, 16, 19] that any fully connected feedforward network \mathcal{N} with ReLU activations partitions the input space into a collection of affine regions, i.e., regions within which \mathcal{N} behaves as an affine map (see Chapter 5). This property has motivated various approaches to studying such networks in recent years. In the special case where the input space is one-dimensional, the affine structure corresponds to a sequence of breakpoints. Telgarsky used this framework to construct a narrow, deep network with an exponential number of affine regions (breakpoints) [15].

However, this breakpoint perspective fails when the input space has dimension greater than one. More recently, such systems have been analyzed from the perspectives of space folds [20, 21, 22] and tropical geometry [12, 18, 23]. The former reveals patterns of self-similarity in the network and emphasizes the emergence of non-convex behavior, while the latter, which we adopt in this paper, facilitates the counting of boundary pieces and affine regions [18] (see Chapter 6 and Chapter 7). It also establishes a connection between ReLU networks and tropical rational maps (see, for example, Proposition 5.1.7), endowing the input space with the structure of a polyhedral complex that can be further studied (see, for example, Chapters 6 and 7).

1.2 Complexity of ReLU Networks

As mentioned above, we measure the complexity of ReLU networks in two ways. First, by counting the number of affine regions partitioning input-space. Second, in the setting of binary classification, by considering the number of linear pieces in the decision boundary.

Especially the former has been studied in recent years. Montafúr et al. [11], building upon work by Pascanu et al. [17], show that a ReLU network $\mathcal{N}: \mathbb{R}^d \rightarrow \mathbb{R}^{n_L}$ with L layers, whose width n_l at layer l satisfies $d \leq n_l \leq w$ for some $w \in \mathbb{N}$, can compute functions that have $\Omega\left(\left(\frac{w}{d}\right)^{(L-1)d} w^d\right)$ affine regions. Using tropical geometry, Zhang et al. [12] derive the upper bound $\mathcal{O}\left(w^{d(L-1)}\right)$ on the maximum number of affine regions (recreating the result of Raghu et al. [13]). Combining these two observations implies that the maximum number of affine regions grows polynomially with the width w and exponentially with the number of layers L . We refer to the regime in which the number of affine regions follows these laws as the *exponential complexity regime*.

However, achieving exponential expressivity typically requires careful network design. Telgarsky [15] constructs a narrow, deep network with exponentially many affine regions. We conjecture that the exponential complexity regime requires such careful constructions, making it unstable and occupy a small volume in parameter-space.

To further quantify the expressive advantage of depth at initialization, Hanin et al. [16] studied the expected number of affine regions in a ReLU network. They found the

complexity to increase linearly in the number of hidden neurons along any one-dimensional subspace, suggesting that, in practice, networks use much less than their theoretically maximal available expressivity. In a subsequent study [14], they show that the expected number of affine regions defined by a deep ReLU network grows polynomially in the number of hidden neurons, with the exponent equal to the input dimension. We call the regime in which the number of affine regions follows these laws the *subexponential complexity regime*. We conjecture that it is more stable than the exponential complexity regime, occupying a larger volume in parameter-space. This stability allows networks to adjust their complexity as needed, helping mitigate overfitting.

The goal of this work is to use tropical geometry to formalize and provide evidence for our claim about the stability of the two complexity regimes. By employing counting techniques and studying upper convex hulls, we aim to show that deep ReLU networks typically do not exploit their theoretically available exponential expressivity. Specifically, in Chapter 9, we introduce a network that we believe resides in the exponential complexity regime. Corollary 9.4.13 demonstrates that the complexity is expected to decrease after applying one random layer, suggesting that the network naturally transitions from the exponential to the subexponential complexity regime.

Chapter 2

Mathematical Background

In this chapter, we introduce key concepts and some background for later parts of this work. Fix an integer $d \in \mathbb{N}$ throughout it.

2.1 Neural Networks

This section covers the basics of fully connected feedforward artificial neural networks. We begin by defining their structure, followed by a brief discussion on training and their applications in classification and regression tasks.

2.1.1 Fully Connected Feedforward Networks

A linear, fully connected feedforward network $\mathcal{N}: \mathbb{R}^d \rightarrow \mathbb{R}^{n_L}$ with L layers can be thought of as a concatenation of L functions $\mathcal{N} = \mathcal{N}_1 \circ \dots \circ \mathcal{N}_L$, where each \mathcal{N}_i acts like $\mathcal{N}_i: \mathbb{R}^{n_{i-1}} \rightarrow \mathbb{R}^{n_i}$ for some natural numbers n_0, \dots, n_L (with $n_0 = d$). Each of the functions \mathcal{N}_i consists of n_i computational units – its coordinate functions, also called *neurons* – which form the i 'th layer of the network. A network is called deep if $L \gg 1$ and wide if $n_i \gg 1$.

In some settings, it may be useful to post-compose \mathcal{N} with the identity map, $\mathcal{N}_0 := \text{id}_{\mathbb{R}^d}$. This *input layer* serves the purpose of feeding the data into the network. The last layer is called the *output layer*. All other layers are called *hidden layers*. We will study them in the following few paragraphs.

Fix $i > 1$. The j 'th neuron in the i 'th layer is assigned a *weight* $\mathbf{w}_{i,j} \in \mathbb{R}^{n_{i-1}}$ and a *bias* $b_{i,j} \in \mathbb{R}$. Given an input $\mathbf{x} \in \mathbb{R}^{n_{i-1}}$, the neuron computes $a_{i,j} := \langle \mathbf{x}, \mathbf{w}_{i,j} \rangle + b_{i,j} \in \mathbb{R}$. The value $a_{i,j}$ is called the *pre-activation* of neuron j in layer i .

To simplify notation, the n_i neurons in layer i are typically grouped into a *weight-matrix* $\mathbf{W}_i \in \mathbb{R}^{n_i, n_{i-1}}$, which contains $\mathbf{w}_{i,j}$ as its j 'th row. Similarly, the biases $b_{i,j}$ are collected into a *bias-vector* $\mathbf{b}_i \in \mathbb{R}^{n_i}$. With this notation, the pre-activation of layer i is given by

$$\mathcal{N}_i(\mathbf{x}) = \mathbf{W}_i \mathbf{x} + \mathbf{b}_i \in \mathbb{R}^{n_i}. \quad (2.1)$$

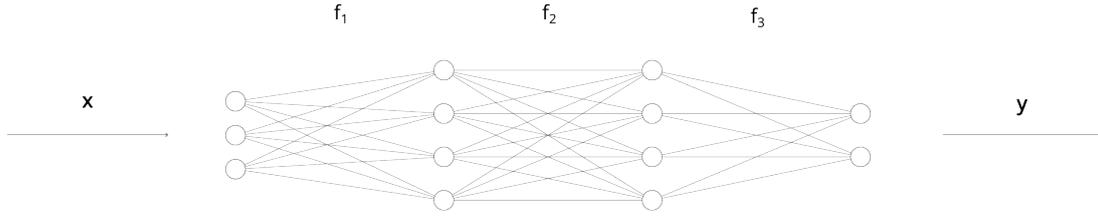


Figure 2.1: The computational graph of a neural network with $L = 3$ layers. The vertical arrays of nodes correspond to the neurons in a layer, with the first (or leftmost) stack resembling the input to the network. The consecutive layers $\mathcal{N}_1, \mathcal{N}_2, \mathcal{N}_3$ propagate the input \mathbf{x} forward through the network, ultimately producing the output $\mathbf{y} := \mathbf{a}_3$.

Given that the neural network is the concatenation of the layer functions \mathcal{N}_i , the recursive definition of the network is

$$\mathbf{a}_0 := \mathbf{x} \tag{2.2}$$

$$\mathbf{a}_{i+1} := f_{i+1}(\mathbf{a}_i) = \mathbf{W}_{i+1}\mathbf{a}_i + \mathbf{b}_{i+1}, \quad i = 0, \dots, L-1, \tag{2.3}$$

where \mathbf{x} is the *input* and $\mathbf{y} := \mathbf{a}_L$ the *output* or *prediction* of the neural network. Equation (2.3) explains how one can think of \mathbf{x} as propagating forward through the network. This process is called the *forward pass*.

Networks of the form given in Equation (2.3) are composed of affine maps and are therefore themselves affine. To enhance their expressivity and expand the class of functions they can model, nonlinearities are typically introduced at every layer. The following definition summarizes the constructions discussed so far.

Definition 2.1.1 (Fully Connected Feedforward Networks). A fully connected feedforward network $\mathcal{N}: \mathbb{R}^d \rightarrow \mathbb{R}^{n_L}$ takes as an input a vector $\mathbf{x} \in \mathbb{R}^d$ and returns an output $\mathbf{y} := \mathbf{a}_L$. It is defined inductively by

$$\begin{cases} \mathbf{a}_0 := \mathbf{x} \\ \mathbf{a}_{l+1} = \rho_{t_{l+1}}(\mathbf{W}_{l+1}\mathbf{a}_l + \mathbf{b}_{l+1}), \quad 0 = 1, \dots, L-1, \end{cases}$$

where $\mathbf{W}_{l+1} \in \mathbb{R}^{n_{l+1}, n_l}$ and $\mathbf{b}_{l+1} \in \mathbb{R}^{n_{l+1}}$ are the *weight matrix* and *bias vector* at layer $l+1$. Furthermore, $\rho_{t_{l+1}}(x) = \max(x, t_{l+1})$ is the *activation function* at layer $l+1$ with *threshold* $t_l \in \mathbb{R} \cup \{-\infty\}$. The number L is called the *depth* of the network, while n_l is the *width* of layer l . The network is *deep* if $L \gg 1$.

Typical nonlinearities at layer l are of the form

$$\begin{aligned} \rho_{t_l}: \mathbb{R} &\rightarrow \mathbb{R} \\ x &\mapsto \max(x, t_l), \end{aligned}$$

where $t_l \in \mathbb{R} \cup \{-\infty\}$ is the *threshold*, and applied element-wise. There are two specific cases relevant to this work:

1. $\rho_0(x) = \max(x, 0)$ is the ReLU (“rectified linear unit”)
2. $\rho_{-\infty}(x) = x$ is the identity.

Definition 2.1.2 (ReLU Networks). A network in the sense of Definition 2.1.1 with ReLU activations (and potentially a linear activation at the last layer) is called a *ReLU network*.

A more comprehensive comparison of different activation functions used in deep learning can be found in [24].

Remark 2.1.3. In this work, we occasionally think of the *activations* \mathbf{a}_l at layer l as a function of the input $\mathbf{x} \in \mathbb{R}^d$, i.e., $\mathbf{a}_l: \mathbb{R}^d \rightarrow \mathbb{R}^{n_l}$.

It is well known [11, 13, 14, 16, 19] that a ReLU network \mathcal{N} partitions the input space into regions on which \mathcal{N} is affine:

Definition 2.1.4 (Affine Regions). Let \mathcal{N} be a ReLU network in the sense of Definition 2.1.1. An *affine region* defined by \mathcal{N} is a maximal connected subset $C \subseteq \mathbb{R}^d$ such that \mathcal{N} restricts to an affine region on C .

Alternatively, the affine regions can be defined as the connected components of the set $\mathbb{R}^d \setminus \{\mathbf{x} \in \mathbb{R}^d \mid \nabla \mathcal{N} \text{ is discontinuous at } \mathbf{x}\}$ [14, Definition 2].

2.1.2 Classification vs. Regression

The fundamental goal of (supervised) machine learning is to learn unknown functions from samples (x_i, y_i) of input-output pairs. Whenever y_i can take continuous values, the task is called *regression*. An example would be inferring the price of a stock from economic information [25]. Whenever y_i can only take discrete values, the task is called *classification*, and y_i is referred to as the *class* or *label* of x_i . An example would be identifying hand-written digits or differentiating images of different objects [26, 27].

While architectures for regression tasks can look as general as the one introduced in Definition 2.1.1, architectures for classification usually require some specifications. In this work we are specifically interested in the case of binary classification, which is the setting where the labels can only take one of two possible values:

Definition 2.1.5 (ReLU Binary Classification Network). A *binary classification network* is a neural network in the sense of Definition 2.1.1, where the last layer has width $n_L = 1$ and is linear (i.e., has threshold $t_L = -\infty$). Throughout this work, we furthermore assume that $t_l = 0$ for $l < L$ and thus speak of a *ReLU (binary) classification network*.

The output of a binary classification network is interpreted as a vote for the class label:

Definition 2.1.6. A *scoring function* takes the output returned by a binary classification network and classifies the corresponding sample. In our setting, it will take the form

$$s: \mathbb{R} \rightarrow \{-1, 1\}$$

$$\tilde{y} \mapsto \text{sign}(y).$$

A binary classification network, together with a scoring function, partitions the input-space into disjoint subsets, each labeled as positive or negative. The boundaries separating positively and negatively labeled subsets are collectively referred to as the *decision boundary* (see Figure 5.1b for an example):

Definition 2.1.7. Let \mathcal{N} be a ReLU binary classification network in the sense of Definition 2.1.5. Then the *decision boundary* of \mathcal{N} is the set

$$\mathcal{B} := \mathcal{N}^{-1}(0). \quad (2.4)$$

2.1.3 Training Neural Networks

The power of artificial neural networks lies in their numerous degrees of freedom and their ability to adjust them in a data-dependent manner. This is achieved by defining a *loss function* \mathcal{L} over a *training set* $(\mathbf{x}_i, \mathbf{y}_i)_{i \in I}$ of samples, which the network aims to minimize. In the setting of regression, for example, this loss might measure how good a prediction \mathbf{a}_L describes a true output \mathbf{y} using the mean squared error:

$$\mathcal{L} = \sum_{i \in I} \frac{1}{|I|} \|\mathbf{a}_L(\mathbf{x}_i) - \mathbf{y}_i\|_2^2.$$

A common method for automatically adjusting the network parameters to minimize the loss \mathcal{L} is *gradient descent*, which iteratively updates the networks parameters like

$$\theta \mapsto \theta - \alpha \nabla_{\theta} \mathcal{L}, \quad (2.5)$$

where $\alpha > 0$ is the *learning rate*. Intuitively, gradient descent adjusts the current estimate of the network parameters by iteratively moving down the loss landscape in the direction of steepest descent. Since it uses the gradient of the loss to guide updates, gradient descent is also called a *first order* optimization technique.

Since the true gradient of the loss function is typically not accessible, it must be estimated empirically using the training data. This leads to a training algorithm called *stochastic gradient descent* (SGD). Note that, in practice, refined versions of SGD are commonly used, such as the ADAM optimizer [28].

In this work, we also examine an optimization algorithm called *Guess & Check* (G&C) [6]. This optimizer operates without gradients and thus falls in the category of *zero'th order* optimization techniques. It works by randomly sampling parameter vectors until it finds one that minimizes the training error (see Algorithm 1). While this algorithm is not commonly used in practice, it serves as a way to study the structure of the loss landscape, as discussed in Chapter 1.

2.2 On Sets, Functions and (In)Dependence

In this section, we present a number of useful definitions and basic mathematical statements.

We start with a number of definitions regarding the relative position of objects in \mathbb{R}^{d+1} .

Algorithm 1 Guess and Check Algorithm for Sampling Parameters θ . Returns the first set of parameters achieving training loss below a tolerance ε .

```

1: Input: Hyperparameter  $\varepsilon$  defining the algorithms tolerance
2: Initialize:  $\mathcal{L} \leftarrow \infty$ 
3: while  $\mathcal{L} \geq \varepsilon$  do
4:   Uniformly sample a random parameter vector  $\theta$ 
5:   Compute the training error  $\mathcal{L}(\theta)$ 
6:   if  $\mathcal{L}(\theta) < \varepsilon$  then
7:     Return  $\theta$ 
8:   end if
9: end while

```

Definition 2.2.1 (Point-Function). Let $\mathcal{N}: \mathbb{R}^d \rightarrow \mathbb{R}$ be a function and $(\mathbf{x}, y) \in \mathbb{R}^{d+1}$ be a point.

- i) We write $(\mathbf{x}, y) \in f$ if (\mathbf{x}, y) lies in the graph of f , i.e., $y = f(\mathbf{x})$.
- ii) We say that (\mathbf{x}, y) lies *above* f if $y > f(\mathbf{x})$. In this case, we write $\mathbf{x} \succ f$. If (\mathbf{x}, y) lies above or on f , i.e., $y \geq f(\mathbf{x})$, we write $\mathbf{x} \succeq f$. Similarly, we write $\mathbf{x} \prec f$ if (\mathbf{x}, y) lies below f and $\mathbf{x} \preceq f$ if (\mathbf{x}, y) lies below or on f .

Definition 2.2.2 (Set-Function). Given a function $\mathcal{N}: \mathbb{R}^d \rightarrow \mathbb{R}$ and a set $X \subseteq \mathbb{R}^{d+1}$, we write $f \succ X$ if $f \succ \mathbf{x}$ for all $\mathbf{x} \in X$. We analogously define $f \succeq X$, $f \prec X$ and $f \preceq X$.

Definition 2.2.3 (Set-Point). Given a point $(\mathbf{x}, y) \in \mathbb{R}^{d+1}$ and a subset $U \subseteq \mathbb{R}^{d+1}$, we say that (\mathbf{x}, y) lies below U if $y < u$ for all $(\mathbf{x}, u) \in U$.

Next, we provide a number of statements regarding linear- and affine independence.

Definition 2.2.4 (Affine Independence). A finite set $\{\mathbf{x}_1, \dots, \mathbf{x}_n\} \subseteq \mathbb{R}^d$ is affinely independent if there does not exist a set of scalars $\{\alpha_1, \dots, \alpha_n\} \subseteq \mathbb{R}$, not all zero, such that

$$\sum_{i=1}^n \alpha_i \mathbf{x}_i = \mathbf{0} \quad \text{and} \quad \sum_{i=1}^n \alpha_i = 0.$$

The following lemma establishes a well-known close relationship between linear and affine independence:

Lemma 2.2.5. *A finite set $\{\mathbf{x}_1, \dots, \mathbf{x}_n\} \subseteq \mathbb{R}^d$ is affinely independent if and only if the set $\{\mathbf{x}_2 - \mathbf{x}_1, \dots, \mathbf{x}_n - \mathbf{x}_1\}$ is linearly independent.*

Proof. “ \Rightarrow ”: Let $\mathbf{x}_1, \dots, \mathbf{x}_n$ be affinely independent and $\alpha_2, \dots, \alpha_n \in \mathbb{R}$ s.t.

$$\sum_{i=2}^n \alpha_i (\mathbf{x}_i - \mathbf{x}_1) = \mathbf{0}.$$

Then, after defining

$$\alpha_1 := - \sum_{i=2}^n \alpha_i,$$

it holds that

$$\sum_{i=1}^n \alpha_i \mathbf{x}_i = 0 \quad \text{and} \quad \sum_{i=1}^n \alpha_i = 0.$$

By affine independence of $\mathbf{x}_1, \dots, \mathbf{x}_n$, this implies that

$$\alpha_i = 0 \quad \forall i = 1, \dots, n.$$

This shows the first implication.

“ \Leftarrow ” Assume that $(\mathbf{x}_i - \mathbf{x}_1)_{i=2, \dots, n}$ are linearly independent and that

$$\sum_{i=1}^n \alpha_i \mathbf{x}_i = 0 \quad \text{with} \quad \sum_{i=1}^n \alpha_i = 0.$$

Then

$$\begin{aligned} 0 &= \sum_{i=1}^n \alpha_i \mathbf{x}_i \\ &= \sum_{i=1}^n \alpha_i \mathbf{x}_1 + \sum_{i=2}^n \alpha_i (\mathbf{x}_i - \mathbf{x}_1) \\ &= \sum_{i=2}^n \alpha_i (\mathbf{x}_i - \mathbf{x}_1). \end{aligned}$$

It follows from linear independence of the differences that $\alpha_i = 0$ for all $i = 2, \dots, n$ and thus also for $i = 1$. This concludes the proof. \square

The relationship between linear and affine independence can be used to determine the affine dimension of a convex hull. However, before proceeding, we first need to define the dimension of a set:

Definition 2.2.6 (Affine Dimension of a Set). Let $X \subseteq \mathbb{R}^d$. Then the *affine hull* of X is the smallest affine subspace of \mathbb{R}^d containing X . We denote it by $\text{affhul}(X)$. We furthermore define the dimension of X as the dimension of its affine hull,

$$\dim X := \dim \text{affhul}(X).$$

Here, the dimension of an affine space is the cardinality of a maximal affinely independent set generating it.

Lemma 2.2.7. Let $S = \{\mathbf{x}_1, \dots, \mathbf{x}_n\} \in \mathbb{R}^d$ be as set of points. Then the convex hull of S has dimension

$$\dim \mathcal{C}(S) = \dim \text{span}(\mathbf{x}_i - \mathbf{x}_1 \mid 2 \leq i \leq n).$$

Proof. Follows from the observation that the smallest affine subspace containing S is the same as the smallest affine subspace containing $\mathcal{C}(S)$ and Lemma 2.2.5. \square

The last statement of this section generalizes the statement of Lemma 2.2.5 to sums of sets:

Lemma 2.2.8. Let $A = \{\mathbf{a}_1, \dots, \mathbf{a}_n\}, B = \{\mathbf{b}_1, \dots, \mathbf{b}_m\} \subseteq \mathbb{R}^d$ be two finite sets of points. Then the set $C := \{a + b \mid a \in A, b \in B\}$ ¹ is affinely independent if and only if the set $\{\mathbf{a}_2 - \mathbf{a}_1, \dots, \mathbf{a}_n - \mathbf{a}_1, \mathbf{b}_2 - \mathbf{b}_1, \dots, \mathbf{b}_m - \mathbf{b}_1\}$ is linearly independent.

Proof. For ease of notation, we introduce the short hand notation $[n] := \{1, \dots, n\}$ for some $n \in \mathbb{N}$. “ \Rightarrow ”: Assume C is affinely independent and let $\alpha_2, \dots, \alpha_n \subseteq \mathbb{R}$ and $\beta_2, \dots, \beta_m \subseteq \mathbb{R}$ be scalars s.t.

$$\sum_{i=2}^n \alpha_i (\mathbf{a}_i - \mathbf{a}_1) + \sum_{j=2}^m \beta_j (\mathbf{b}_j - \mathbf{b}_1) = 0.$$

After defining

$$\alpha_1 := - \sum_{i=2}^n \alpha_i$$

$$\beta_1 := - \sum_{j=2}^m \beta_j,$$

this statement can be rephrased as

$$\sum_{i=1}^n \alpha_i \mathbf{a}_i + \sum_{j=1}^m \beta_j \mathbf{b}_j = 0 \tag{2.6}$$

with

$$\sum_{i=1}^n \alpha_i = 0 \tag{2.7}$$

and

$$\sum_{j=1}^m \beta_j = 0. \tag{2.8}$$

Next, we define the scalars $\gamma_{i,j} \in \mathbb{R}$ for $1 \leq i \leq n, 1 \leq j \leq m$ as

$$\gamma_{i,j} := \frac{\alpha_i}{m} + \frac{\beta_j}{n}.$$

Then, by Equations (2.7)-(2.8),

$$\sum_{i=1}^n \gamma_{i,j} = \beta_j \tag{2.9}$$

$$\sum_{j=1}^m \gamma_{i,j} = \alpha_i. \tag{2.10}$$

¹We will later recognize this sum as the Minkowski sum of A and B

Consequently, Equation (2.6) can be re-written as

$$\sum_{i,j \in [n] \times [m]} \gamma_{i,j}(\mathbf{a}_i + \mathbf{b}_j) = 0$$

with

$$\sum_{i,j \in [n] \times [m]} \gamma_{i,j} = 0$$

(the latter equality follows from Equations (2.7)-(2.8)).

By assumption, C is affinely independent and we conclude that $\gamma_{i,j} = 0$ for all $i, j \in [n] \times [m]$. It follows from Equations (2.9)-(2.10) that $\alpha_i = 0$ and $\beta_j = 0$ for all $1 \leq i \leq n$, $1 \leq j \leq m$. This shows one direction.

“ \Leftarrow ”: Assume the set $\{\mathbf{a}_2 - \mathbf{a}_1, \dots, \mathbf{a}_n - \mathbf{a}_1, \mathbf{b}_2 - \mathbf{b}_1, \dots, \mathbf{b}_m - \mathbf{b}_1\}$ is linearly independent and let $\gamma_{i,j} \in \mathbb{R}$ be scalars for all $i, j \in [n] \times [m]$ s.t.

$$\sum_{i,j \in [n] \times [m]} \gamma_{i,j}(\mathbf{a}_i + \mathbf{b}_j) = 0 \quad \text{and} \quad \sum_{i,j \in [n] \times [m]} \gamma_{i,j} = 0. \quad (2.11)$$

Define

$$\alpha_i := \sum_{j=1}^m \gamma_{i,j}$$

$$\beta_j := \sum_{i=1}^n \gamma_{i,j}.$$

Then Equation (2.11) can be re-written as

$$\sum_{i=1}^n \alpha_i \mathbf{a}_i + \sum_{j=1}^m \beta_j \mathbf{b}_j = 0 \quad \text{and} \quad \sum_{i=1}^n \alpha_i 0 + \sum_{j=1}^m \beta_j 0 = 0.$$

Consequently,

$$\begin{aligned} \sum_{i=1}^n \alpha_i \mathbf{a}_i + \sum_{j=1}^m \beta_j \mathbf{b}_j &= \sum_{i=1}^n \alpha_i \mathbf{a}_1 + \sum_{j=1}^m \beta_j \mathbf{b}_1 + \sum_{i=2}^n \alpha_i (\mathbf{a}_i - \mathbf{a}_1) + \sum_{j=2}^m \beta_j (\mathbf{b}_j - \mathbf{b}_1) \\ &= \sum_{i=2}^n \alpha_i (\mathbf{a}_i - \mathbf{a}_1) + \sum_{j=2}^m \beta_j (\mathbf{b}_j - \mathbf{b}_1) \\ &= 0. \end{aligned}$$

By assumption, we conclude that $\alpha_i = 0$ and $\beta_j = 0$ for all $2 \leq i \leq n$, $2 \leq j \leq m$. This also implies $\alpha_1 = \beta_1 = 0$ and thus $\gamma_{i,j} = 0$ for all $i, j \in [n] \times [m]$. This shows the other direction and concludes the proof. \square

2.3 Polyhedral Complexes

In this section, we introduce fundamental concepts related to polyhedral complexes, which will play a crucial role in later chapters.

Definition 2.3.1 (Polyhedron & Polytope). A *polyhedron* P is the intersection of finitely many closed half-spaces in \mathbb{R}^d and can thus be written as

$$P = \{\mathbf{x} \in \mathbb{R}^d \mid \mathbf{A}\mathbf{x} \leq \mathbf{b}\}$$

for some $\mathbf{A} \in \mathbb{R}^{m,d}$ and $\mathbf{b} \in \mathbb{R}^m$ with $m \in \mathbb{N}$. A *polytope* is a bounded polyhedron.

As the following lemma shows, this is not an empty definition.

Lemma 2.3.2 ([29, Section 2.3]). *The convex hull of finitely many vertices is a polytope.*

The “boundary” of a polyhedron is made up of faces:

Definition 2.3.3 (Face of a Polyhedron). Let $P \subseteq \mathbb{R}^d$ be a polyhedron. A *face* of P is a non-empty subset $F \subseteq P$ s.t.

$$F = P \cap \{\mathbf{x} \mid \mathbf{A}'\mathbf{x} = \mathbf{b}'\}, \quad (2.12)$$

where \mathbf{A}' arises from \mathbf{A} and \mathbf{b}' arises from \mathbf{b} by deleting rows with the same indices. In other words, some of the inequalities defining P are satisfied as equalities in F (see also Definition 2.3.7). The *dimension* of a face is the dimension of the smallest affine subspace containing it. Zero-dimensional faces are called *vertices*.

Useful for our studies will be a collection of faces called the upper convex hull:

Definition 2.3.4 (Upper Convex Hull). Consider the polytope P formed by the convex hull $\mathcal{C}(S)$ of a finite set of points $S \subseteq \mathbb{R}^d$. An *upper face* of P is a face whose inner normal vector (the normal vector pointing inward to P) has negative last coordinate. We call the union of all upper faces the *upper convex hull* of S , denoted by $\mathcal{U}(S)$. The union of all k -faces in $\mathcal{U}(S)$ is denoted by $\mathcal{U}_k(S)$. In the specific case where $k = 0$, we write $\mathcal{U}^*(S) := \mathcal{U}_0(S)$ (see Figure 3.3 for an example).

This concludes our study of individual polyhedra. A collection of multiple polyhedra can form a structure known as a polyhedral complex:

Definition 2.3.5 (Polyhedral Complex). A *polyhedral complex* Σ is a collection of polyhedra satisfying two conditions:

- i) if P is a polyhedron contained in Σ , then any face of P is also contained in Σ ,
- ii) if P and Q are both polyhedra contained in Σ , then $P \cap Q$ is either empty or a face of both P and Q .

The *support* of a polyhedral complex $\Sigma \in \mathbb{R}^d$ is the union of the points in all of the polyhedra:

$$|\Sigma| := \{\mathbf{x} \in \mathbb{R}^d \mid \mathbf{x} \in P \text{ for some polyhedron } P \in \Sigma\}.$$

Individual polyhedra inside Σ are called *cells*. The *dimension* of a cell σ is the dimension of the smallest affine subspace containing it. The $(d - 1)$ -dimensional polyhedra are called *facets* of Σ .

A polyhedral complex consists of skeletons of different dimensionalities:

Definition 2.3.6 (K-Skeleton). The k -skeleton of a polyhedral complex Σ is the sub-complex formed by all cells $\sigma \in \Sigma$ with dimension $\dim(\sigma) = k$.

This concludes our introduction to polyhedral complexes. In the remainder of this section, we develop a theory for determining the dimension of a polytope, largely following the approach outlined in [30].

For simplicity, we will use the short-hand notation $\{\mathbf{Ax} \geq \mathbf{b}\} := \{\mathbf{x} \mid \mathbf{Ax} \geq \mathbf{b}\}$.

Definition 2.3.7 (Implicit Equalities). Given a system of linear inequalities $\mathbf{Ax} \geq \mathbf{b}$ (that is, $\langle \mathbf{a}_i, \mathbf{x} \rangle \geq b_i$ for all rows \mathbf{a}_i in \mathbf{A} and entries b_i in \mathbf{b}), an inequality $\langle \mathbf{a}_i, \mathbf{x} \rangle \geq b_i$ is an *implicit equality* if

$$\langle \mathbf{a}_i, \mathbf{x} \rangle = b_i \quad \forall \mathbf{x} \in \{\mathbf{Ax} \geq \mathbf{b}\}.$$

We define $\mathbf{A}_=\mathbf{x} = \mathbf{b}_=$ to be the system of implicit equalities and $\mathbf{A}_+\mathbf{x} \geq \mathbf{b}_+$ to be the system of remaining inequalities.

Intuitively, the implicit inequalities can be thought of as restricting the polyhedron to a lower-dimensional subspace. For example, a two-dimensional polyhedron embedded in three-dimensional space needs to be confined to a two-dimensional affine subspace by an implicit equality constraint.

The following is an almost trivial helping lemma:

Lemma 2.3.8 ([30, Proposition 8]). Let $P = \{\mathbf{Ax} \geq \mathbf{b}\}$ be a polyhedron such that not all of the constraints are implicit equality constraints. Then there exists a point $\mathbf{x} \in P$ such that

$$\begin{aligned} \mathbf{A}_=\mathbf{x} &= \mathbf{b}_= \\ \mathbf{A}_+\mathbf{x} &> \mathbf{b}_+. \end{aligned}$$

In other words, there exists a point in P that does not lie on a face of P .

Proof. Let $\mathbf{x} \in P$. If \mathbf{x} has the desired properties, we are done. If not, there exists an inequality constraint which is satisfied with equality by \mathbf{x} (that is, \mathbf{x} lies on a face of P). Since this constraint is not part of the equality constraints, there must be exists an $\mathbf{x}_0 \in P$ satisfying the constraint with strict inequality. Pick this \mathbf{x}_0 and repeat the argument. \square

The previous lemma allows identifying the affine hull of a polyhedron:

Lemma 2.3.9 ([30, Lemma 9]). Let $P = \{\mathbf{Ax} \geq \mathbf{b}\}$ be a polyhedron with existing implicit equality constraints. Then

$$\text{affhul}(P) = \{\mathbf{A}_=\mathbf{x} = \mathbf{b}_=\}.$$

Proof. If all of the constraints are implicit equality constraints, the claim is trivial. Hence, we may assume that there exist some constraints that are not implicit equality constraints.

" \subseteq ": We start by showing that

$$\text{affhul}(P) \subseteq \{\mathbf{A}_=\mathbf{x} = \mathbf{b}_=\}. \tag{2.13}$$

Let $\mathbf{x} \in \text{affhul}(P)$. Then there exists an $m \in \mathbb{N}$ and $\mathbf{x}_1, \dots, \mathbf{x}_m \in P$, $\alpha_1, \dots, \alpha_m \in \mathbb{R}$ satisfying $\sum_{i=1}^m \alpha_i = 1$ such that

$$\mathbf{x} = \sum_{i=1}^m \alpha_i \mathbf{x}_i.$$

This implies that

$$\begin{aligned} \mathbf{A}_=\mathbf{x} &= \sum_{i=1}^m \alpha_i \mathbf{A}_=\mathbf{x}_i \\ &= \sum_{i=1}^m \alpha_i \mathbf{b}_= \\ &= \mathbf{b}_=, \end{aligned}$$

where we used the fact that $P \subseteq \{\mathbf{A}_=\mathbf{x} = \mathbf{b}_=\}$. This shows the first inclusion.

" \supseteq ": Next we show that

$$\text{affhul}(P) \supseteq \{\mathbf{A}_=\mathbf{x} = \mathbf{b}_=\}. \quad (2.14)$$

Let \mathbf{x} satisfy $\mathbf{A}_=\mathbf{x} = \mathbf{b}_=$ and pick an $\mathbf{x}' \in P$ s.t.

$$\begin{aligned} \mathbf{A}_=\mathbf{x}' &= \mathbf{b}_= \\ \mathbf{A}_+\mathbf{x}' &> \mathbf{b}_+, \end{aligned}$$

which is possible by Lemma 2.3.8.

If $\mathbf{x} \in P$, we are done. Also if $\mathbf{x} = \mathbf{x}'$, then $\mathbf{x} \in P$ and we are done. Hence, we may assume that $\mathbf{x} \neq \mathbf{x}'$ and $\mathbf{x} \notin P$. Let l be the line-segment from \mathbf{x}' to \mathbf{x} , i.e., $l = \{t\mathbf{x} + (1-t)\mathbf{x}' \mid t \in [0, 1]\}$. We claim that

$$|l \cap P| > 1, \quad (2.15)$$

where $|\cdot|$ denotes the cardinality of a set. Since $\mathbf{x}' \in l \cap P$, we only need to find one more point in the intersection (see Figure 2.2). Since $\mathbf{A}_=\mathbf{x} = \mathbf{b}_=$ and $\mathbf{x} \notin P$, we know that

$$\mathbf{A}_+\mathbf{x} \not\geq \mathbf{b}_+.$$

Additionally using the fact that

$$\begin{aligned} \mathbf{A}_=\mathbf{y} &= \mathbf{b}_= \quad \forall \mathbf{y} \in l \\ \mathbf{A}_+\mathbf{x}' &> \mathbf{b}_+ \end{aligned}$$

we conclude there has to exist an $\mathbf{x}_0 \in l$, $\mathbf{x}_0 \neq \mathbf{x}'$, satisfying

$$\begin{aligned} \mathbf{A}_=\mathbf{x}_0 &= \mathbf{b}_= \\ \mathbf{A}_+\mathbf{x}_0 &\geq \mathbf{b}_+ \end{aligned}$$

and thus $\mathbf{x}_0 \in P \cap l$. This shows Equation (2.15).

Finally, since $\mathbf{x}', \mathbf{x}_0 \in P$ and $\mathbf{x} \in l$,

$$\text{affhul } P \supseteq \text{affhul}\{\mathbf{x}_0, \mathbf{x}'\} \ni \mathbf{x}.$$

This concludes the proof. □

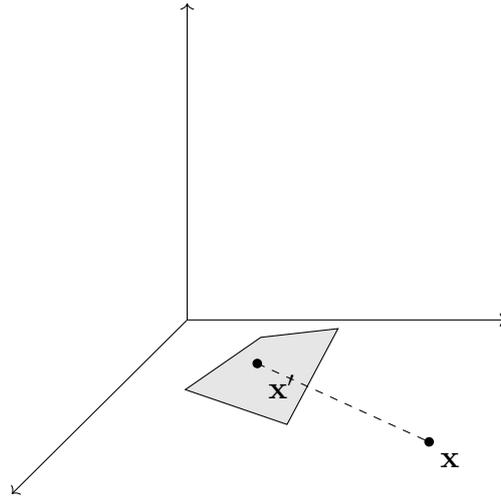


Figure 2.2: Inspired by Figure 5.5 in [30]. Line segment joining x' and x .

Proposition 2.3.10 ([30, Corollary 10]). *Let $P = \{\mathbf{Ax} \geq \mathbf{b}\}$ be a non-empty polyhedron with existing implicit equality constraints. Then*

$$\dim P = d - \text{rank } \mathbf{A}_= . \quad (2.16)$$

Proof. By Definition 2.2.6, $\dim P = \dim \text{affhul}(P)$. The proposition then follows from Lemma 2.3.9 and the rank-nullity-theorem. \square

Chapter 3

Affine Geometry

In this chapter, we introduce fundamental concepts of affine geometry, covering basic definitions, the dual representation of affine functions, and their connection to upper convex hulls. We also explore tessellations induced by maxima over affine functions.

Throughout this chapter, fix an integer $d \in \mathbb{N}$.

3.1 Affine and (D)CPA Functions

We begin by introducing fundamental concepts.

Definition 3.1.1 (Affine Functions). Given a vector $\mathbf{a} \in \mathbb{R}^d$ and a scalar $b \in \mathbb{R}$, we define the affine function with parameters \mathbf{a} and b as

$$\begin{aligned} f_{\mathbf{a},b}: \mathbb{R}^d &\rightarrow \mathbb{R} \\ \mathbf{x} &\mapsto \langle \mathbf{a}, \mathbf{x} \rangle + b, \end{aligned}$$

where $\langle \cdot, \cdot \rangle$ is the standard Euclidean inner product on \mathbb{R}^d .

Ultimately, we will be taking maxima over affine maps. To classify such functions, we introduce the following concept:

Definition 3.1.2 (CPA Functions). We say that a function $f: \mathbb{R}^d \rightarrow \mathbb{R}$ is CPA if it is convex and piecewise affine. We denote by $\text{CPA}(d)$ that set of CPA functions $\mathbb{R}^d \rightarrow \mathbb{R}$.

It turns out that the class of CPA functions coincides with the class of maxima over affine functions:

Proposition 3.1.3 (Characterizing CPA Functions [18, Proposition 2]). *Any function $F: \mathbb{R}^d \rightarrow \mathbb{R}$ of the form*

$$F(\mathbf{x}) = \max\{f_1(\mathbf{x}), \dots, f_n(\mathbf{x})\}$$

with affine functions $f_i: \mathbb{R}^d \rightarrow \mathbb{R}$ is CPA. Also every CPA function with a finite number of affine pieces is of this form.

Later in this work, we will also consider differences of CPA functions:

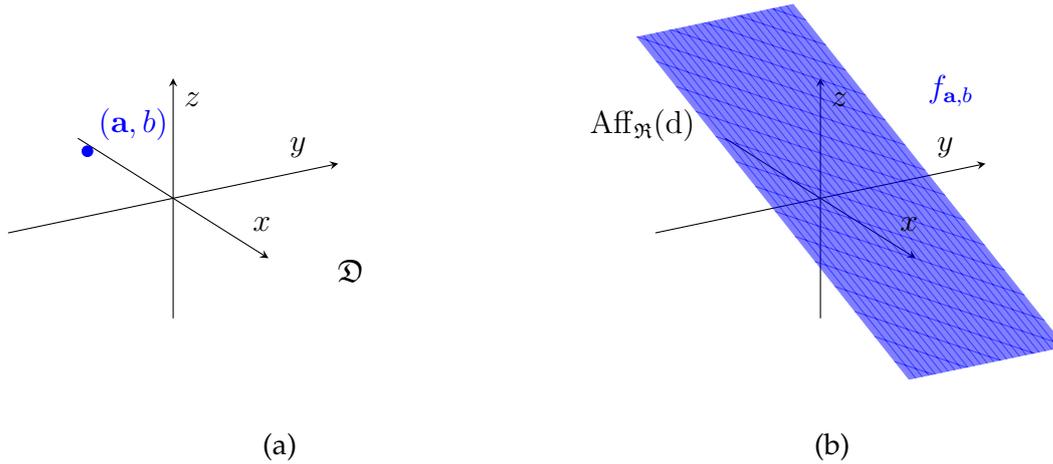


Figure 3.1: Example of the dual representation of an affine map $f_{\mathbf{a},b}$ with $\mathbf{a} = (-1/2, -3/4)$, $b = 3/4$. Subfigure (b) contains the graph of $f_{\mathbf{a},b} \in \text{Aff}_{\mathfrak{R}}(d)$ and Subfigure (a) contains the parameterizing dual point $(\mathbf{a}, b) \in \mathcal{D}$. The map \mathcal{R} assigns to the point (\mathbf{a}, b) the affine map $f_{\mathbf{a},b}$.

Definition 3.1.4 (DCPA Functions). We say that a function $f: \mathbb{R}^d \rightarrow \mathbb{R}$ is DCPA if it can be written as the difference of two CPA functions. We denote by $\text{DCPA}(d)$ the set of DCPA function $\mathbb{R}^d \rightarrow \mathbb{R}$.

3.2 Affine Dualities

In this section, we mainly follow the construction presented in [18], which allows mapping an affine function $f: \mathbb{R}^d \rightarrow \mathbb{R}$ to a “dual space”. As an outlook, exploring this transformation will ultimately lead to understanding how ReLU networks can be understood as DCPA functions.

The graph of an affine function $\mathbb{R}^d \rightarrow \mathbb{R}$ defines a hyperplane in *real space*, which we define as $\mathfrak{R} := \mathbb{R}^d \times \mathbb{R} = \mathbb{R}^{d+1}$. The space of affine functions whose graph lies in \mathfrak{R} is called *real affine space*, denoted by $\text{Aff}_{\mathfrak{R}}(d)$.

As mentioned in Definition 3.1.1, any affine function $f_{\mathbf{a},b} \in \text{Aff}_{\mathfrak{R}}(d)$ is characterized by its parameters $(\mathbf{a}, b) \in \mathbb{R}^{d+1}$. We refer to the copy of \mathbb{R}^{d+1} that parametrizes affine functions in $\text{Aff}_{\mathfrak{R}}(d)$ as the *dual space* \mathcal{D} .

The following lemma is a natural consequence of this construction, as it allows translating between real affine space and dual space:

Lemma 3.2.1. *For any fixed dimension d , there exists a bijection between dual space and real affine space, given by*

$$\begin{aligned} \mathcal{R}: \mathcal{D} &\xrightarrow{\sim} \text{Aff}_{\mathfrak{R}}(d) \\ (\mathbf{x}, y) &\mapsto f_{\mathbf{x},y}. \end{aligned}$$

An example for \mathcal{R} can be found in Figure 3.1. It has the following properties.

Proposition 3.2.2. Let $\{\mathbf{x}_i, y_i\}_{i=1, \dots, n} \subseteq \mathfrak{D}$ be a set of dual points. Then the following are true:

i) \mathcal{R} is a linear operator, i.e., for any set of scalars $\{\alpha_i\}_{i=1, \dots, n} \subseteq \mathbb{R}$,

$$\mathcal{R} \left(\sum_{i=1}^n \alpha_i (\mathbf{x}_i, y_i) \right) = \sum_{i=1}^n \alpha_i \mathcal{R}((\mathbf{x}_i, y_i)).$$

ii) The set of dual points is linearly independent if and only if the corresponding set $\{\mathcal{R}((\mathbf{x}_i, y_i))\}_{i=1, \dots, n}$ of affine functions is linearly independent.

iii) The set of dual points is affinely independent if and only if the corresponding set $\{\mathcal{R}((\mathbf{x}_i, y_i))\}_{i=1, \dots, n}$ of affine functions is affinely independent.

Proof. i) can be confirmed by an easy calculation. ii) follows from i) and iii) from ii) and Lemma 2.2.5. \square

Since both \mathfrak{R} and \mathfrak{D} are copies of \mathbb{R}^{d+1} , it is natural to ask whether we can reverse their roles in the above construction. The answer to this question is yes. We define *dual affine space* $\text{Aff}_{\mathfrak{D}}(d)$ as the space of affine functions with graph in \mathfrak{D} . Analogously to the above construction, these affine functions are parameterized by points in \mathfrak{R} , though with a slight caveat:

Lemma 3.2.3. For any fixed dimension d , there exists a bijection between dual affine space and real space. It is given by

$$\begin{aligned} \tilde{\mathcal{R}}: \text{Aff}_{\mathfrak{D}}(d) &\xrightarrow{\sim} \mathfrak{R} \\ f_{\mathbf{a}, b} &\mapsto (-\mathbf{a}, b). \end{aligned}$$

Figure 3.2 provides an overview the relationship between \mathfrak{R} , \mathfrak{D} , $\text{Aff}_{\mathfrak{R}}(d)$ and $\text{Aff}_{\mathfrak{D}}(d)$.

$$\begin{array}{ccc} \text{Aff}_{\mathfrak{R}}(d) & \xrightarrow{\text{graph}} & \mathfrak{R} \\ \uparrow \mathcal{R} & & \uparrow \tilde{\mathcal{R}} \\ \mathfrak{D} & \xleftarrow{\text{graph}} & \text{Aff}_{\mathfrak{D}}(d) \end{array}$$

Figure 3.2: Diagram indicating the relationship between real (affine) and dual (affine) space.

Note that, compared to \mathcal{R} , the function $\tilde{\mathcal{R}}$ includes an additional minus and maps in the opposite direction. This is essential for ensuring that the duality properties in the following proposition hold:

Proposition 3.2.4 (Duality Properties [18, Proposition 7]). *The maps \mathcal{R} and $\tilde{\mathcal{R}}$ have the following properties (using notation from Definition 2.2.2):*

1. A dual point $\mathbf{c} \in \mathfrak{D}$ lies on the graph of a dual affine function $f_{\mathbf{a},\mathbf{b}} \in \text{Aff}_{\mathfrak{D}}(\mathfrak{d})$ if and only if the graph of the corresponding real affine function $\mathcal{R}(\mathbf{c})$ contains the corresponding real point $\check{\mathcal{R}}(f_{\mathbf{a},\mathbf{b}})$:

$$\mathbf{c} \in f_{\mathbf{a},\mathbf{b}} \iff \check{\mathcal{R}}(f_{\mathbf{a},\mathbf{b}}) \in \mathcal{R}(\mathbf{c})$$

2. A dual point $\mathbf{c} \in \mathfrak{D}$ lies above the graph of a dual affine function $f_{\mathbf{a},\mathbf{b}} \in \text{Aff}_{\mathfrak{D}}(\mathfrak{d})$ if and only if the real point $\check{\mathcal{R}}(f_{\mathbf{a},\mathbf{b}})$ lies below the graph of $\mathcal{R}(\mathbf{c})$:

$$\mathbf{c} \succ f_{\mathbf{a},\mathbf{b}} \iff \mathcal{R}(\mathbf{c}) \succ \check{\mathcal{R}}(f_{\mathbf{a},\mathbf{b}})$$

3.3 CPA Functions as Upper Convex Hulls

In the previous section, we explored a duality that enables us to identify affine maps with the vector containing their parameters. In this section, we apply these results to maxima over affine functions, which, by Proposition 3.1.3, can be understood as CPA functions.

In light of the duality results from the previous section, CPA functions correspond to finite sets of dual points:

Definition 3.3.1. On the set $\mathcal{P}_{\text{fin}}(\mathfrak{D})$ of finite subsets of \mathfrak{D} , the operator

$$\begin{aligned} \mathcal{Q}: \mathcal{P}_{\text{fin}}(\mathfrak{D}) &\rightarrow \text{CPA}(\mathfrak{d}) \\ S &\mapsto \mathcal{Q}(S) := \max_{\mathbf{s} \in S} \mathcal{R}(\mathbf{s}) \end{aligned}$$

assigns to a set of dual points the associated CPA function

$$\max_{\mathbf{s} \in S} \mathcal{R}(\mathbf{s})(\mathbf{x}) = \max_{(\mathbf{a},b) \in S} \langle \mathbf{x}, \mathbf{a} \rangle + b.$$

We define $\mathcal{Q}(\emptyset) := 0$. On a vector of finite sets of dual points, \mathcal{Q} acts component-wise.

Note that, by Proposition 3.1.3, the operator \mathcal{Q} does indeed map to $\text{CPA}(\mathfrak{d})$.

Our next objective is to establish a connection between CPA functions and upper convex hulls. To begin, we first state the following proposition:

Proposition 3.3.2 (Maximality of Upper Convex Hull [18, Proposition 9]). *Let $S \subseteq \mathfrak{D}$ be a finite set of points. Then for every point $w \in \mathfrak{D}$ lying below or on $\mathcal{U}(S)$ (in the sense of Definition 2.2.3), the affine function dual to w lies fully below the maximum of the affine functions whose duals lie in $\mathcal{U}^*(S)$. That is,*

$$\mathcal{R}(w) \leq \max\{\mathcal{R}(s) \mid s \in \mathcal{U}^*(S)\} = \mathcal{Q}(\mathcal{U}^*(S)). \quad (3.1)$$

If w lies truly below $\mathcal{U}(S)$, then even

$$\mathcal{R}(w) < \mathcal{Q}(\mathcal{U}^*(S)). \quad (3.2)$$

Proof. The proof follows largely the same structure as [18, Proposition 9], with a few minor adaptations.

Let $(\mathbf{x}_1, y_1), \dots, (\mathbf{x}_n, y_n) \in \mathfrak{D}$, $n \geq 3$, be distinct dual points. We start with the following two observations:

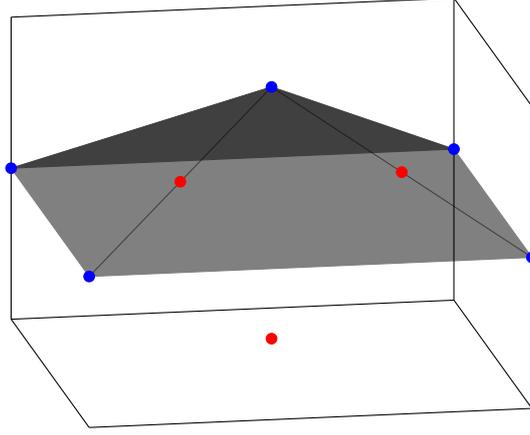


Figure 3.3: Example of an upper convex hull. Let S be the union of all displayed points. The blue points correspond to $\mathcal{U}^*(S)$, the black surface is $\mathcal{U}_2(S)$. In particular, $\mathcal{Q}(S)$ is uniquely identified by only the blue points.

- i) if (\mathbf{x}_1, y_1) lies directly below (\mathbf{x}_2, y_2) , i.e., $\mathbf{x}_1 = \mathbf{x}_2$ and $y_1 < y_2$, then the dual plane related to (\mathbf{x}_1, y_1) lies below (\mathbf{x}_2, y_2) , i.e., $\mathcal{R}((\mathbf{x}_1, y_1))(\mathbf{x}) < \mathcal{R}((\mathbf{x}_2, y_2))(\mathbf{x})$ for all $\mathbf{x} \in \mathbb{R}^d$
- ii) if (\mathbf{x}_n, y_n) lies on a face of $\mathcal{U}(S)$ spanned by $(\mathbf{x}_1, y_1), \dots, (\mathbf{x}_{n-1}, y_{n-1}) \in \mathcal{U}^*(S)$, then $\mathcal{R}(\mathbf{x}_n, y_n) \leq \max\{\mathcal{R}((\mathbf{x}_i, y_i)) \mid i = 1, \dots, n-1\}$.

Claim *i*) is trivial. For claim *ii*), assume there exist $\alpha_i \in [0, 1]$, $\sum_{i=1}^n \alpha_i = 1$, s.t.

$$(\mathbf{x}_n, y_n) = \sum_{i=1}^{n-1} \alpha_i (\mathbf{x}_i, y_i).$$

Then

$$\mathcal{R}((\mathbf{x}_n, y_n))(\mathbf{x}) = \sum_{i=1}^{n-1} \alpha_i \mathcal{R}((\mathbf{x}_i, y_i))(\mathbf{x}) \quad \forall \mathbf{x} \in \mathbb{R}^d$$

by linearity of \mathcal{R} (see Proposition 3.2.2). In particular,

$$\mathcal{R}(\mathbf{x}_n, y_n)(\mathbf{x}) \leq \max\{\mathcal{R}((\mathbf{x}_i, y_i))(\mathbf{x}) \mid i = 1, \dots, n-1\} \quad \forall \mathbf{x} \in \mathbb{R}^d.$$

This shows claim *ii*).

The proposition then follows from the following observation. Assume that the point w lies below or on $\mathcal{U}(S)$. Let (\mathbf{x}_1, y_1) be a point directly above w lying on $\mathcal{U}(S)$. Then, by *i*), $\mathcal{R}(w) < \mathcal{R}((\mathbf{x}_1, y_1))$ if w does not lie on $\mathcal{U}(S)$ and $\mathcal{R}(w) \leq \mathcal{R}((\mathbf{x}_1, y_1))$ otherwise. Furthermore, by *ii*), $\mathcal{R}((\mathbf{x}_1, y_1)) \leq \max\{\mathcal{R}(s) \mid s \in \mathcal{U}^*(S)\}$. This shows the claim. \square

Having established this proposition, the identification of CPA functions with upper convex hulls is a corollary:

Corollary 3.3.3 (CPAs as Upper Convex Hulls). *Every CPA function $\mathcal{Q}(S)$ can be uniquely represented as an upper convex hull in dual space. That is, $\mathcal{Q}(S) = \mathcal{Q}(\mathcal{U}^*(S))$*

Proof. Let $Q(S)$ be a CPA function. Then for any $\mathbf{x} \in \mathbb{R}^d$,

$$\begin{aligned} Q(S)(\mathbf{x}) &= \max_{s \in S} \mathcal{R}(s)(\mathbf{x}) \\ &= \max\{\max_{s \in \mathcal{U}^*(S)} R(s)(\mathbf{x}), \max_{s \in S \setminus \mathcal{U}^*(S)} \mathcal{R}(s)(\mathbf{x})\} \\ &\stackrel{3.3.2}{=} \max_{s \in \mathcal{U}^*(S)} R(s)(\mathbf{x}) \\ &= Q(\mathcal{U}^*(S))(\mathbf{x}). \end{aligned}$$

This shows the claim. □

A visualization of Corollary 3.3.3 can be found in Figure 3.3.

3.4 Tessellations

CPA functions induce a *tessellation* of \mathbb{R}^d . It plays an important role in understanding ReLU networks:

Definition 3.4.1 (Tessellation). Given a CPA function $F(\mathbf{x}) := \max\{f_1(\mathbf{x}), \dots, f_n(\mathbf{x})\}$, a *cell* induced by F is

$$\{\mathbf{x} \in \mathbb{R}^d \mid f_i(\mathbf{x}) = f_{i'}(\mathbf{x}) \geq f_j(\mathbf{x}) \text{ for all } i, i' \in I, j \in J\},$$

where I, J are disjoint sets whose union is $\{1, 2, \dots, n\}$. The set of all cells induced F is called the *tessellation* induced by F and denoted by $\mathcal{T}(F)$.

Figure 5.1a contains an example of a tessellation. By a slight abuse of notation, we will write $\mathcal{T}(S)$ for the tessellation induced by the CPA function $Q(S)$.

The following lemma establishes a connection between tessellations and polyhedral complexes, which were discussed in Section 2.3:

Lemma 3.4.2. *The tessellation induced by a CPA function F forms a polyhedral complex.*

Proof. Every cell of F is a polyhedron since it is defined by a set of linear inequalities. It is left to show that the following two properties hold (see Definition 2.3.5):

- i) any face of an cell is also a cell,
- ii) the intersection of two cells is either empty or a face of both cells.

But this follows directly from the definition of the tessellation. Indeed, let σ be a cell defined by two sets I and J , as in Definition 3.4.1. Then a face of σ is a cell associated with two sets $I' \supseteq I, J' \subseteq J$ obtained by moving indices from J to I (note that, at a face of σ , there are more active equality constraints). This shows *i*). To see that *ii*) holds, observe that the intersection of two cells, associated with sets I, J and I', J' , respectively, is the cell associated with the sets $I \cap I', J \cup J' \cup (I \Delta I')$ (here, Δ denotes the symmetric difference of sets). □

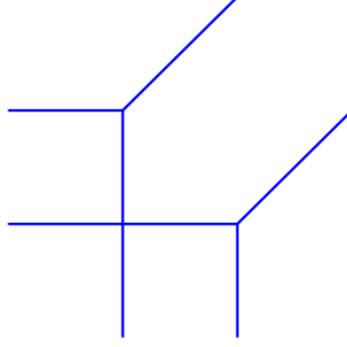


Figure 3.4: Figure 1 in [12]. Example of a tessellation, induced by the DCPA-function given in Equation (3.3).

This last lemma tells us that we may think of a tessellation as a polyhedral complex and the following definition makes sense:

Definition 3.4.3. Let F be a CPA function. We denote by $\mathcal{T}_k(F)$ the k -skeleton of the tessellation induced by F . The support of $\mathcal{T}_{d-1}(F)$ is also called an *affine (or tropical) hypersurface*.

Example 3.4.4. As an example, Figure 3.4 shows the tessellation induced by P and N .

$$f: \mathbb{R}^2 \rightarrow \mathbb{R}, \quad (x, y) \mapsto \max\{1 + 2x, 1 + 2y, 2 + x + y, 2 + x, 2 + y, 2\}. \quad (3.3)$$

The blue lines correspond to points on which two affine functions agree and are larger than the others. They form the 1-skeleton of the tessellation. The intersections of these lines are the 0-cells. On each of the white convex regions (the 2-cells) f is affine.

Figure 3.4 also illustrates how the tessellation forms a polyhedral complex: the face of any polyhedron is again a polyhedron (for example, the faces of the white convex regions are the 1-cells), and the intersection of any two polyhedra is either empty or again a face.

So far, we have studied what it means for a CPA function to induce a tessellation of \mathbb{R}^d . The following definition clarifies what it means for a DCPA to do so.

Definition 3.4.5. Let $F = Q(P) - Q(N)$ be a DCPA function. We then define the tessellation $\mathcal{T}(P, N)$ induced by F to consist of all non-empty pairwise intersections of cells induced by P and N , i.e.

$$\mathcal{T}(P, N) := \{\sigma \cap \sigma' \mid \sigma \in \mathcal{T}(P), \sigma' \in \mathcal{T}(N)\} \setminus \emptyset.$$

As it turns out, $\mathcal{T}(P, N)$ is closely related to tessellations induced by different CPA functions:

Definition 3.4.6 (Refinements). Let \mathcal{T} and \mathcal{F} be tessellations of \mathbb{R}^d . We say that \mathcal{T} is a *refinement* \mathcal{F} if every cell of \mathcal{T} is contained in a cell of \mathcal{F} . In this case, we write $\mathcal{T} \ll \mathcal{F}$.

Lemma 3.4.7. *Given two sets of dual points $P, N \subseteq \mathcal{D}$, it holds that*

$$\mathcal{T}(P \cup N) \ll \mathcal{T}(P, N) \ll \mathcal{T}(N).$$

Proof. For ease of notation, enumerate $N = \{n_1, \dots, n_m\}$ and $P = \{p_1, \dots, p_k\}$ with $m, k \in \mathbb{N}$.

A cell of $\mathcal{T}(N)$ is given by the solution of a system

$$\begin{cases} \mathcal{R}(n_i) = \mathcal{R}(n_j) & \forall i, j \in I \\ \mathcal{R}(n_i) \geq \mathcal{R}(n_j) & \forall i \in I, j \in J \end{cases} \quad (3.4)$$

for some disjoint partition $I \sqcup J = \{1, \dots, m\}$.

A cell of $\mathcal{T}(P, N)$ is given by the solution of a system

$$\begin{cases} \mathcal{R}(n_i) = \mathcal{R}(n_j) & \forall i, j \in I \\ \mathcal{R}(p_{i'}) = \mathcal{R}(p_{j'}) & \forall i', j' \in I' \\ \mathcal{R}(n_i) \geq \mathcal{R}(n_j) & \forall i \in I, j \in J \\ \mathcal{R}(p_{i'}) \geq \mathcal{R}(p_{j'}) & \forall i' \in I', j' \in J' \end{cases} \quad (3.5)$$

for some disjoint partitions $I \sqcup J = \{1, \dots, m\}$ and $I' \sqcup J' = \{1, \dots, k\}$.

A cell of $\mathcal{T}(P \cup N)$ is given by the solution of a system

$$\begin{cases} \mathcal{R}(n_i) = \mathcal{R}(n_j) & \forall i, j \in I \\ \mathcal{R}(p_{i'}) = \mathcal{R}(p_{j'}) & \forall i', j' \in I' \\ \mathcal{R}(n_i) = \mathcal{R}(p_{j'}) & \forall i \in I, j' \in I' \\ \mathcal{R}(n_i) \geq \mathcal{R}(n_j) & \forall i \in I, j \in J \\ \mathcal{R}(p_{i'}) \geq \mathcal{R}(p_{j'}) & \forall i' \in I', j' \in J' \\ \mathcal{R}(n_i) \geq \mathcal{R}(p_{j'}) & \forall i \in I, j' \in J' \\ \mathcal{R}(p_{i'}) \geq \mathcal{R}(n_j) & \forall i' \in I', j \in J \end{cases} \quad (3.6)$$

for some disjoint partitions $I \sqcup J = \{1, \dots, m\}$ and $I' \sqcup J' = \{1, \dots, k\}$.

Clearly, any solution to System (3.6) also solves System (3.5) and any solution to System (3.5) also solves System (3.4). This implies the claim. \square

Chapter 4

Tropical Geometry

There exists a second perspective on affine geometry, which we introduced in the previous chapter, called tropical geometry. This chapter aims to lay out the foundational concepts of this more abstract framework, as it provides an alternative viewpoint on the constructions presented in the following chapters. It allows us to equip the space of affine functions (with integer weights and allowing for $-\infty$ biases) with the structure of a semi-ring. Our approach in this chapter largely follows [12] and [29].

A reader solely interested in deriving the dual representation of ReLU networks may skip this chapter and still understand the subsequent results. However, we believe that the tropical perspective offers a valuable alternative viewpoint which helps develop a deeper understanding of the underlying geometry.

4.1 Basic Definitions

Tropical geometry takes place in the *tropical semiring*:

Definition 4.1.1. The *tropical semiring* consists of the set $\mathbb{T} := \mathbb{R} \cup \{-\infty\}$ together with the operations \oplus and \odot , where \oplus is called *tropical addition* and \odot is called *tropical multiplication*. These operations are defined as $x \oplus y := \max\{x, y\}$ and $x \odot y := x + y$ for $x, y \in \mathbb{R}$. Furthermore, $-\infty \oplus x := x$ and $-\infty \odot x := -\infty$. The tropical quotient is defined as $x \oslash y := x - y$.

Remark 4.1.2. As the name suggests, $(\mathbb{T}, \oplus, \odot)$ is a semi-ring. For completeness, we recall the defining properties below:

- i) (\mathbb{R}, \odot) is a monoid under tropical multiplication, i.e., \odot is associative and has a multiplicative identity: 0
- ii) (\mathbb{T}, \oplus) is an abelian group *except for the existence of a tropical additive inverse*, i.e., \oplus is associative, commutative and has an additive identity: $-\infty$
- iii) tropical multiplication is distributive with respect to tropical addition.

Next, we introduce the notion of tropical exponentiation:

Definition 4.1.3. Given $x \in \mathbb{T}$ and $a \in \mathbb{N}$, $x^{\odot a}$ is defined as

$$x^{\odot a} := \begin{cases} x \odot \dots \odot x = a \cdot x & x \in \mathbb{R} \\ -\infty & x = -\infty \text{ and } a > 0. \\ 0 & x = -\infty \text{ and } a = 0 \end{cases}$$

Similar to the real case, we can define monomials/polynomials over \mathbb{T} :

Definition 4.1.4 (Tropical Monomials). A *tropical monomial* in d variables is an expression of the form

$$b \odot x_1^{\odot a_1} \odot \dots \odot x_d^{\odot a_d}$$

where $b \in \mathbb{T}$ and $a_1, \dots, a_d \in \mathbb{N}$. Multiindex notation allows writing $b \odot x^{\odot \alpha}$ where $\alpha \in \mathbb{N}^d$. We denote the space of tropical monomials in d variables by $\mathbb{T}\{x_1, \dots, x_d\}$.

Definition 4.1.5 (Tropical Polynomials). A *tropical polynomial* f is a finite tropical sum of tropical monomials,

$$f = b_1 \odot x^{\odot \alpha_1} \oplus \dots \oplus b_r \odot x^{\odot \alpha_r}$$

where $\alpha_i \in \mathbb{N}^d$ and $b_i \in \mathbb{T}$ for all $i = 1, \dots, r$. We will assume that $\alpha_i \neq \alpha_j$ for $i \neq j$, i.e., that the polynomial is in some sense reduced maximally. We denote the space of tropical polynomials in d variables by $\mathbb{T}[x_1, \dots, x_d]$

A concept that frequently arises in the context of neural networks is that of a *tropical rational*:

Definition 4.1.6 (Tropical Rationals). A *tropical rational* is the tropical quotient of two tropical polynomials f and g :

$$f \oslash g = f - g.$$

We will denote the space of tropical rationals in d variables by $\mathbb{T}(x_1, \dots, x_d)$.

So far, we have discussed concepts from tropical algebra, where the exponents of tropical polynomials are natural numbers. However, this assumption can be relaxed to allow for arbitrary real exponents, though at the expense of losing the structure of tropical algebra. This broader perspective will yield more general results once parallels to ReLU networks are explored in later chapters.

Start by extending Definition 4.1.3 naturally to make sense of what it means to raise an expression to a real-valued tropical power, namely $x^{\odot a} = ax$ for $a \in \mathbb{R}$. This allows for the following definition:

Definition 4.1.7 (Tropical Simple Signomial). A *tropical simple signomial* in d variables is an expression of the form

$$b \odot x_1^{\odot a_1} \odot \dots \odot x_d^{\odot a_d}$$

where $b \in \mathbb{T}$ and $a_1, \dots, a_d \in \mathbb{R}$. In particular, the exponent can be any real number, as opposed to tropical monomials. Multiindex notation allows writing $b \odot x^{\odot \alpha}$ with $\alpha \in \mathbb{R}^d$. We denote the space of tropical simple signomials by $\mathbb{T}_{\mathbb{R}}\{x_1, \dots, x_d\}$.

Analogously to how tropical polynomials are tropical sums of tropical monomials, tropical signomials are tropical sums of tropical simple signomials:

Definition 4.1.8 (Tropical Signomial). A *tropical signomial* φ is a finite tropical sum of tropical monomials

$$\varphi = b_1 \odot x^{\odot \alpha_1} \oplus \dots \oplus b_r \odot x^{\odot \alpha_r}$$

where $\alpha_i \in \mathbb{R}^d$ and $b_i \in \mathbb{T}$ for all $i = 1, \dots, r$. We denote the space of tropical signomials in d variables by $\mathbb{T}_{\mathbb{R}}[x_1, \dots, x_d]$.

Finally, we define tropical quotients of tropical signomials:

Definition 4.1.9 (Tropical Rational Signomial). A *tropical rational signomial* is the tropical quotient of two tropical signomials φ and ψ :

$$\varphi \oslash \psi = \varphi - \psi.$$

We denote the space of tropical rational signomials in d variables by $\mathbb{T}_{\mathbb{R}}(x_1, \dots, x_d)$

4.2 Relation to Affine Geometry

As mentioned in the introduction to this section, there exists a close relationship between tropical geometry and affine geometry. In this section, we formalize and establish this connection.

Affine geometry, as introduced in Chapter 3, does not account for $-\infty$. In order to establish the aforementioned connection, we begin by restricting tropical signomials to those with finite coefficients:

Definition 4.2.1. We denote by

$$\mathbb{T}_{\mathbb{R}}^{\text{fin}}\{x_1, \dots, x_d\} := \{b \odot x^{\odot \alpha} \in \mathbb{T}_{\mathbb{R}}\{x_1, \dots, x_d\} \mid b \in \mathbb{R}\}$$

the space of tropical simple signomials with finite coefficients. Analogously, we denote by $\mathbb{T}_{\mathbb{R}}^{\text{fin}}[x_1, \dots, x_d]$ the space of tropical sums of tropical simple signomials with finite coefficients and by $\mathbb{T}_{\mathbb{R}}^{\text{fin}}(x_1, \dots, x_d)$ the space of tropical quotients of tropical signomials with finite coefficients.

It is evident from this definition that, for instance, $\mathbb{T}_{\mathbb{R}}^{\text{fin}}\{x_1, \dots, x_d\} \subsetneq \mathbb{T}_{\mathbb{R}}\{x_1, \dots, x_d\}$, since the space on the right hand side also permits coefficients $b = -\infty$. Nevertheless, this restriction allows making the following identifications between objects from affine geometry and objects from tropical geometry:

Proposition 4.2.2. *The following maps are bijections.*

1. *Affine functions can be identified with finite-coefficient simple signomials,*

$$\begin{aligned} \text{Aff}_{\mathbb{R}}(d) &\xrightarrow{\sim} \mathbb{T}_{\mathbb{R}}^{\text{fin}}\{x_1, \dots, x_d\} \\ f_{\mathbf{a}, b} &\mapsto b \odot x^{\odot \mathbf{a}} \end{aligned}$$

2. *CPA functions can be identified with finite-coefficient tropical signomials,*

$$\begin{aligned} \text{CPA}(d) &\xrightarrow{\sim} \mathbb{T}_{\mathbb{R}}^{\text{fin}}[x_1, \dots, x_n] \\ \max_{i=1, \dots, n} f_{\mathbf{a}_i, b_i} &\mapsto b_1 \odot x^{\odot \mathbf{a}_1} \oplus \dots \oplus b_n \odot x^{\odot \mathbf{a}_n} \end{aligned}$$

3. DCPA functions can be identified with finite-coefficient tropical rational signomials,

$$\begin{aligned} \text{DCPA}(d) &\xrightarrow{\sim} \mathbb{T}_{\mathbb{R}}^{\text{fin}}(x_1, \dots, x_n) \\ \max_{i=1, \dots, n} f_{\mathbf{a}_i, b_i} - \max_{i=1, \dots, m} f_{\mathbf{c}_i, d_i} &\mapsto \bigoplus_{i=1}^n b_i \odot x^{\odot \mathbf{a}_i} - \bigoplus_{i=1}^m d_i \odot x^{\odot \mathbf{c}_i}. \end{aligned}$$

In the remainder of this work, we will use the identifications in Proposition 4.2.2 without always mentioning them specifically. For example, we will treat tropical signomials $b \odot x^{\odot \mathbf{a}}$ as affine functions

$$\begin{aligned} b \odot x^{\odot \mathbf{a}}: \mathbb{R}^d &\rightarrow \mathbb{R} \\ \mathbf{x} &\mapsto f_{\mathbf{a}, b}(\mathbf{x}). \end{aligned}$$

Chapter 5

Dual Representation of Neural Networks

The previous two chapters introduced affine/tropical geometry. We will now apply this knowledge to establish a connection between fully connected feedforward networks and DCPA functions, allowing us to translate the network to dual space.

5.1 Neural Networks and Affine Geometry

This section covers the affine perspective. As described in Section 4.2, the following results can directly be translated to the tropical setting.

In order to establish the connection, we first need to develop some more machinery, beginning with the definition of how to sum two sets:

Definition 5.1.1 (Minkowski Sum). Given two non-empty sets $X, Y \subseteq \mathbb{R}^{d+1}$, we define

$$X \oplus Y := \{\mathbf{x} + \mathbf{y} \mid \mathbf{x} \in X, \mathbf{y} \in Y\}$$

to be the *Minkowski sum* of X and Y . We define $X \oplus \emptyset := X$. On vectors of sets of dual points, we define $+$ to act component-wise.

Next, we list a number of properties of the operator \mathcal{Q} from Definition 3.3.1, which assigns to a set of dual points the corresponding CPA function:

Lemma 5.1.2 (Properties of \mathcal{Q} [18, Proposition 12]). *For any two sets of points $X, Y \subseteq \mathcal{D}$ and every non-negative scalar $\alpha \geq 0$, the following are true:*

- i) $\mathcal{Q}(X \cup Y) = \max\{\mathcal{Q}(X), \mathcal{Q}(Y)\}$
- ii) $\mathcal{Q}(X \oplus Y) = \mathcal{Q}(X) + \mathcal{Q}(Y)$
- iii) $\alpha \cdot \mathcal{Q} = \mathcal{Q}(\alpha \cdot X)$, where the multiplication on the right hand side is the natural multiplication of a set with a real number.

Proof. Follows directly from the definition of \mathcal{Q} and a simple computation. □

Neural networks rely heavily on matrix multiplications. On our mission to translate them to dual space, we must first establish the concept of matrix multiplication in the dual setting:

Definition 5.1.3. We define the multiplication of an $m \times n$ matrix \mathbf{A} with a vector X of n finite sets of dual points as

$$\begin{aligned} \cdot: \mathbb{R}^{m,n} \times (P_{\text{fin}}(\mathcal{D}))^n &\rightarrow (P_{\text{fin}}(\mathcal{D}))^m \\ (\mathbf{A}, X) &\mapsto \mathbf{A} \cdot X \end{aligned}$$

where

$$(\mathbf{A} \cdot X)_i := \bigoplus_{j=1}^n \mathbf{A}_{ij} \cdot X_j.$$

In the notation above, $(P_{\text{fin}}(\mathcal{D}))^n$ denotes the n -fold Cartesian product of $P_{\text{fin}}(\mathcal{D})$ with itself and $\bigoplus_{j=1}^n$ denotes the Minkowski sum over the sets indexed by $\{1, \dots, n\}$.

The following lemma shows that matrix multiplication and the \mathcal{Q} -operator commute:

Lemma 5.1.4 (Matrix Multiplication [18]). *Let $X \in (P_{\text{fin}}(\mathcal{D}))^n$ be a vector of finite sets of dual points and $\mathbf{A} \in \mathbb{R}_+^{m,n}$ a matrix with non-negative entries. Then*

$$\mathbf{A}\mathcal{Q}(X) = \mathcal{Q}(\mathbf{A} \cdot X).$$

Proof. In the setting of this lemma, \mathcal{Q} is applied coordinate-wise, and thus we can understand the multiplication $\mathbf{A}\mathcal{Q}(X)$ as the matrix-vector multiplication:

$$\begin{aligned} [\mathbf{A}\mathcal{Q}(X)]_i &= \sum_{j=1}^n \mathbf{A}_{ij} [\mathcal{Q}(X)]_j \\ &\stackrel{(*)}{=} \sum_{j=1}^n \mathcal{Q}(\mathbf{A}_{ij} X_j) \\ &\stackrel{(**)}{=} \mathcal{Q}\left(\bigoplus_{j=1}^n \mathbf{A}_{ij} X_j\right) \\ &\stackrel{5.1.3}{=} \mathcal{Q}([\mathbf{A} \cdot X]_i) \\ &= [\mathcal{Q}(\mathbf{A} \cdot X)]_i. \end{aligned}$$

In $(*)$ and $(**)$ we used Lemma 5.1.2 *iii)* and *ii)*, respectively. □

In order to account for biases, we define how to add a scalar to a set of dual points:

Definition 5.1.5. A scalar can be added to a set of dual points by adding the scalar to the last entry of each point in the set:

$$\begin{aligned} \boxplus: P_{\text{fin}}(\mathcal{D}) \times \mathbb{R} &\rightarrow P_{\text{fin}}(\mathcal{D}) \\ (X, \alpha) &\mapsto X \boxplus \alpha, \end{aligned}$$

where $X \boxplus \alpha$ is the set

$$X \boxplus \alpha := \{(\mathbf{x}, y + \alpha) \mid (\mathbf{x}, y) \in X\}.$$

It turns out that \mathcal{Q} is well behaved with respect to \boxplus :

Lemma 5.1.6. *For any finite set of dual points $X \subseteq \mathfrak{D}$ and scalar $\alpha \in \mathbb{R}$, it holds that*

$$\mathcal{Q}(X) + \alpha = \mathcal{Q}(X \boxplus \alpha).$$

Proof. For every $\mathbf{z} \in \mathbb{R}^d$, it holds that

$$\begin{aligned} (\mathcal{Q}(X) + \alpha)(\mathbf{z}) &= \mathcal{Q}(X)(\mathbf{z}) + \alpha \\ &= \max_{(\mathbf{a}, b) \in X} f_{\mathbf{a}, b}(\mathbf{z}) + \alpha \\ &= \max_{(\mathbf{a}, b) \in X} \langle \mathbf{a}, \mathbf{z} \rangle + b + \alpha \\ &= \max_{(\mathbf{a}, b) \in X \boxplus \alpha} \langle \mathbf{a}, \mathbf{z} \rangle + b \\ &= \mathcal{Q}(X \boxplus \alpha)(\mathbf{z}). \end{aligned}$$

□

We are now ready to present the following fundamental proposition that establishes the connection between ReLU networks and differences of piecewise affine functions:

Proposition 5.1.7 (Dual Representation [18, Proposition 16]). *Assume that a neural network in the sense of Definition 2.1.1 can, up to layer $l - 1$, be written as a DCPA function $\mathbf{a}_{l-1} = \mathcal{Q}(P_{l-1}) - \mathcal{Q}(N_{l-1})$ for some vectors of finite sets of dual points P_{l-1}, N_{l-1} . Then, after writing $\mathbf{W}_l = \mathbf{W}_l^+ - \mathbf{W}_l^-$ using matrices \mathbf{W}_l^+ and \mathbf{W}_l^- with non-negative entries, also the network up to the l 'th layer can be written as a DCPA function*

$$\mathbf{a}_l = \mathcal{Q}(P_l) - \mathcal{Q}(N_l)$$

with

$$\begin{aligned} N_l &= (\mathbf{W}_l^- \cdot P_{l-1}) \boxplus (\mathbf{W}_l^+ \cdot N_{l-1}) \\ P_l &= (((\mathbf{W}_l^+ \cdot P_{l-1}) \boxplus (\mathbf{W}_l^- \cdot N_{l-1})) \boxplus \mathbf{b}_l) \cup \begin{cases} N_l \boxplus t_l, & t_l \neq -\infty \\ \emptyset, & t_l = -\infty. \end{cases} \end{aligned}$$

Proof. First, note that

$$\begin{aligned} \mathbf{W}_l \mathbf{a}_{l-1} &= (\mathbf{W}_l^+ - \mathbf{W}_l^-) (\mathcal{Q}(P_{l-1}) - \mathcal{Q}(N_{l-1})) \\ &= (\mathbf{W}_l^+ \mathcal{Q}(P_{l-1}) + \mathbf{W}_l^- \mathcal{Q}(N_{l-1})) - (\mathbf{W}_l^- \mathcal{Q}(P_{l-1}) + \mathbf{W}_l^+ \mathcal{Q}(N_{l-1})) \\ &\stackrel{(*)}{=} \mathcal{Q}((\mathbf{W}_l^+ \cdot P_{l-1}) \boxplus (\mathbf{W}_l^- \cdot N_{l-1})) - \mathcal{Q}((\mathbf{W}_l^- \cdot P_{l-1}) \boxplus (\mathbf{W}_l^+ \cdot N_{l-1})) \\ &= \mathcal{Q}((\mathbf{W}_l^+ \cdot P_{l-1}) \boxplus (\mathbf{W}_l^- \cdot N_{l-1})) - \mathcal{Q}(N_l), \end{aligned}$$

where in (*) we used Lemma 5.1.4 and Lemma 5.1.2 ii).

Next, assume first that $t_l \neq -\infty$. Using the identity $\max\{x - y, t\} = \max\{x, y + t\} - y$ and the above reformulation of $\mathbf{W}_l \mathbf{a}_{l-1}$, we can write, using the definitions of N_l and

P_l claimed in the proposition,

$$\begin{aligned}
\mathbf{a}_l &= \rho(\mathbf{W}_l a_{l-1} + \mathbf{b}_l) \\
&= \max\{\mathcal{Q}((\mathbf{W}_l^+ \cdot P_{l-1}) \diamond (\mathbf{W}_l^- \cdot N_{l-1})) - \mathcal{Q}(N_l) + \mathbf{b}_l, t_l\} \\
&\stackrel{(**)}{=} \max\{\mathcal{Q}(((\mathbf{W}_l^+ \cdot P_{l-1}) \diamond (\mathbf{W}_l^- \cdot N_{l-1})) \boxplus \mathbf{b}_l) - \mathcal{Q}(N_l), t_l\} \\
&= \max\{\mathcal{Q}(((\mathbf{W}_l^+ \cdot P_{l-1}) \diamond (\mathbf{W}_l^- \cdot N_{l-1})) \boxplus \mathbf{b}_l), \mathcal{Q}(N_l) + t_l\} - \mathcal{Q}(N_l) \\
&\stackrel{(***)}{=} \mathcal{Q}((((\mathbf{W}_l^+ \cdot P_{l-1}) \diamond (\mathbf{W}_l^- \cdot N_{l-1})) \boxplus \mathbf{b}_l) \cup (N_l \boxplus t_l)) - \mathcal{Q}(N_l) \\
&= \mathcal{Q}(P_l) - \mathcal{Q}(N_l).
\end{aligned}$$

In (***) we used a vectorized version of Lemma 5.1.6 and in (***) we used Lemma 5.1.2 *i*) and Lemma 5.1.6.

Now, assume $t_l = -\infty$. Then

$$\begin{aligned}
\mathbf{a}_l &= \rho(\mathbf{W}_l a_{l-1} + \mathbf{b}_l) \\
&= \max\{\mathcal{Q}((\mathbf{W}_l^+ \cdot P_{l-1}) \diamond (\mathbf{W}_l^- \cdot N_{l-1})) - \mathcal{Q}(N_l) + \mathbf{b}_l, -\infty\} \\
&= \mathcal{Q}(((\mathbf{W}_l^+ \cdot P_{l-1}) \diamond (\mathbf{W}_l^- \cdot N_{l-1})) \boxplus \mathbf{b}_l) - \mathcal{Q}(N_l) \\
&= \mathcal{Q}(P_l) - \mathcal{Q}(N_l).
\end{aligned}$$

This concludes the proof. □

The following corollary makes sure Proposition 5.1.7 can actually be applied to ReLU networks by establishing the base case:

Corollary 5.1.8. *Every neural network \mathcal{N} in the sense of Definition 2.1.1 can be written as a DCPA function*

$$\mathcal{N} = \mathcal{Q}(P) - \mathcal{Q}(N)$$

for some vectors of dual sets of dual points $P, N \subseteq \mathfrak{D}$. We call (P, N) the dual representation of \mathcal{N} .

Proof. The proof goes by induction on the number of layers L of the network. Crucially, for this proof we introduce an identity-layer $\mathcal{N}_0: \mathbb{R}^d \rightarrow \mathbb{R}^d$ at the beginning of the network (which can be thought of as an input-layer, see Section 2.1). This convenient trick simplifies the proof in the following way:

We start the induction with $L = 0$. In this case, the network just consists of the input-function $\mathcal{N}_0 = \text{id}_{\mathbb{R}^d}$, whose i 'th coordinate function can be represented as

$$\begin{aligned}
\mathbf{x}_i &= \langle \mathbf{x}, \mathbf{e}_i \rangle + 0 = f_{0,i}(\mathbf{x}) \\
&= \mathcal{Q}(\{(\mathbf{e}_i, 0)\})(\mathbf{x}) - \mathcal{Q}(\emptyset)(\mathbf{x}) \\
&= \mathcal{Q}(\{(\mathbf{e}_i, 0)\})(\mathbf{x}),
\end{aligned}$$

where $\mathbf{e}_i \in \mathbb{R}^d$ is the i 'th unit vector with entries $(\mathbf{e}_i)_j = \delta_{ij}$. Hence, we can write

$$\mathcal{N}_0 = \mathcal{Q}(\{(\mathbf{e}_1, 0)\}, \dots, \{(\mathbf{e}_d, 0)\})$$

and may choose $P_0 = (\{(\mathbf{e}_1, 0)\}, \dots, \{(\mathbf{e}_d, 0)\})$, $N_0 = (\emptyset)$. This shows the claim for $L = 0$. For the induction step, use Proposition 5.1.7. □

Remark 5.1.9. Proposition 5.1.7 and Corollary 5.1.8 are important tools for the rest of this work. We want to use this remark to highlight their significance. Let $\mathcal{N}: \mathbb{R}^d \rightarrow \mathbb{R}$ be a ReLU network with L layers.

- i) Given the weights of \mathcal{N} , the dual representation provides a symbolic representation (P_l, N_l) of \mathcal{N} up to layer l .
- ii) It is given by two n_l -dimensional vectors P_l, N_l of finite sets of dual points, where n_l is the width of layer l . The dual points are $d + 1$ -dimensional.
- iii) As we will see later, this symbolic representation allows counting the number of affine regions defined by \mathcal{N} (see Chapter 7). In the case of binary classification, it furthermore allows counting the linear pieces in the decision boundary (see Chapter 6).
- iv) After each layer l , the sets of dual points can be replaced by their upper convex hull (see Corollary 3.3.3). In particular, for every $i \in \{1, \dots, n_l\}$, the set $(P_l)_i \subseteq \mathfrak{D} = \mathbb{R}^{d+1}$ can be replaced by its upper convex hull vertices $\mathcal{U}^*((P_l)_i)$. The same holds for $(N_l)_i$.

Scattered throughout this work, we employ a running example to highlight the above points. See Example 5.2.6, Example 6.2.5 and Example 7.1.12.

5.1.1 Positive and Negative Samples

In this subsection, we use the duality result from Proposition 5.1.7 to understand whether a point $\mathbf{x} \in \mathbb{R}^d$ is positively or negatively classified by $\mathcal{N} = \mathcal{Q}(P) - \mathcal{Q}(N)$.

We start with the following lemma:

Lemma 5.1.10. *Let $A, B \subseteq \mathbb{R}^{d+1}$ be finite sets of points. Then the following are true for $\mathbf{x} \in \mathbb{R}^d$:*

$$\mathcal{Q}(A \cup B)(\mathbf{x}) = \mathcal{Q}(A)(\mathbf{x}) \iff \mathcal{Q}(A)(\mathbf{x}) \geq \mathcal{Q}(B)(\mathbf{x}) \quad (5.1)$$

$$\mathcal{Q}(A \cup B)(\mathbf{x}) = \mathcal{Q}(B)(\mathbf{x}) \iff \mathcal{Q}(A)(\mathbf{x}) \leq \mathcal{Q}(B)(\mathbf{x}). \quad (5.2)$$

Proof. Follows from Lemma 5.1.2 i). □

This lemma allows identifying positively and negatively labeled samples:

Proposition 5.1.11. *Let $\mathcal{N} = \mathcal{Q}(P) - \mathcal{Q}(N)$ be a ReLU binary classification network. Then the following are true for an input $\mathbf{x} \in \mathbb{R}^d$:*

$$\mathcal{N}(\mathbf{x}) \geq 0 \iff \mathcal{Q}(P \cup N)(\mathbf{x}) = \mathcal{Q}(P)(\mathbf{x}) \quad (5.3)$$

$$\mathcal{N}(\mathbf{x}) \leq 0 \iff \mathcal{Q}(P \cup N)(\mathbf{x}) = \mathcal{Q}(N)(\mathbf{x}). \quad (5.4)$$

Proof. Follows from Lemma 3.4.7 and Lemma 5.1.10. □

In words, Proposition 5.1.11 states that positively labeled points are characterized by the property that the maximum $\mathcal{Q}(P \cup N)$ is attained by $\mathcal{Q}(P)$. Similarly, negatively labeled points are characterized by the property that the maximum is attained by $\mathcal{Q}(N)$.

5.1.2 Counting Dual Points

In this section, we use the recursive formulation in Proposition 5.1.7 to work towards a better understanding of the sets inside P and Q . Specifically, we are interested in upper bounding the size of the sets inside P_l and N_l at layer l . We start with a basic definition.

Definition 5.1.12. Given a vector of finite sets of dual points $X \in \mathcal{P}_{\text{fin}}(\mathcal{D})^n$, we define $s(X) := \max_{i=1, \dots, n} |X_i|$ to be the size of the largest set in X .

In the following lemma, we establish some properties of the function s .

Lemma 5.1.13. Let $\mathbf{A} \in \mathbb{R}^{m,n}$ be a matrix, $b \in \mathbb{R}$ a scalar and $X \in (\mathcal{P}_{\text{fin}}(\mathcal{D}) \setminus \emptyset)^n$, $Y \in (\mathcal{P}_{\text{fin}}(\mathcal{D}) \setminus \emptyset)^k$ be vectors of finite, non-empty sets of dual points. Then the following are true:

i) For any $1 \leq i \leq m$,

$$|(\mathbf{A} \cdot X)_i| \leq \prod_{j=1}^n |X_j|.$$

In particular,

$$s(\mathbf{A} \cdot X) \leq s(X)^n.$$

ii)

$$|X_i \diamond Y_j| \leq |X_i| |Y_j| \quad \forall 1 \leq i \leq n, 1 \leq j \leq k.$$

iii)

$$|X_i \boxplus b| = |X_i| \quad \forall 1 \leq i \leq n.$$

Proof. ii) and iii) are clear. For i) we compute

$$\begin{aligned} |(\mathbf{A} \cdot X)_i| &= \left| \bigoplus_{j=1}^n \mathbf{A}_{ij} X_j \right| \\ &\stackrel{ii)}{\leq} \prod_{j=1}^n |\mathbf{A}_{ij} X_j| \\ &\leq \prod_{j=1}^n |X_j| \\ &\leq s(X)^n. \end{aligned}$$

□

To better understand the maximal size of the sets in P_l and N_l , we first define the quantity

$$\xi_l := s(P_l) s(N_l), \quad \text{for } l = 1, \dots, L.$$

The following lemma provides recursive bounds involving ξ_l , $s(N_l)$ and $s(P_l)$:

Lemma 5.1.14. For any $l = 2, \dots, L$, the quantities $s(N_l)$ and $s(P_l)$ can be upper-bounded in terms of ξ_{l-1} via

$$s(N_l) \leq \xi_{l-1}^{n_{l-1}} \quad (5.5)$$

and

$$s(P_l) \leq \begin{cases} 2\xi_{l-1}^{n_{l-1}}, & t_k \neq -\infty \\ \xi_{l-1}^{n_{l-1}}, & t_k = -\infty. \end{cases} \quad (5.6)$$

Furthermore, ξ_l satisfies the recursive relation

$$\xi_l \leq \begin{cases} 2\xi_{l-1}^{2n_{l-1}}, & t_k \neq -\infty \\ \xi_{l-1}^{2n_{l-1}}, & t_k = -\infty. \end{cases} \quad (5.7)$$

Proof. We start by showing Equation (5.5). Fix $l \in \{2, \dots, L\}$. Then

$$\begin{aligned} |(N_l)_i| &= |(\mathbf{W}_l^- \cdot P_{l-1})_i \diamond (\mathbf{W}_l^+ \cdot N_{l-1})_i| \\ &\stackrel{5.1.13}{\leq} |(\mathbf{W}_l^- \cdot P_{l-1})_i| |(\mathbf{W}_l^+ \cdot N_{l-1})_i| \\ &\stackrel{5.1.13}{\leq} s(P_{l-1})^{n_{l-1}} s(N_{l-1})^{n_{l-1}} \\ &= \xi_{l-1}^{n_{l-1}} \end{aligned}$$

This shows Equation (5.5).

Similarly, one can show that

$$s(P_{l-1}) \leq s(P_{l-1})^{n_{l-1}} s(N_{l-1})^{n_{l-1}} + \begin{cases} s(P_{l-1})^{n_{l-1}} s(N_{l-1})^{n_{l-1}}, & t_{l-1} \neq \infty \\ 0, & t_{l-1} = \infty. \end{cases}$$

which shows Equation (5.6). Finally, Equation (5.7) follows from the other two.. \square

Lemma 5.1.14 provides recursive bounds for $s(N_l)$ and $s(P_l)$ involving ξ_l . In order to derive closed-form bounds for $s(N_L)$ and $s(P_L)$, we need to unravel these recursions. This is accomplished in the following proposition:

Proposition 5.1.15. Consider the setting of Proposition 5.1.7 with initial definitions of N_0 and P_0 as in Corollary 5.1.8. Then, if the thresholds t_l satisfy $t_l \neq -\infty$ for all $l \leq k$, the product ξ_k is upper bounded by

$$\xi_k \leq 2^{1+\sum_{j=1}^{k-1} \prod_{i=j}^{k-1} 2^{n_j}}, \quad k = 1, \dots, L. \quad (5.8)$$

Proof. We prove the claim by induction on k .

First note that $N_0 = (\emptyset)$ and $P_0 = (\{(e_1, 0)\}, \dots, \{(e_d, 0)\})$, which immediately implies that $s(P_0) = 1$. It follows directly from Proposition 5.1.7, Lemma 5.1.13.i) and the fact that $\mathbf{W}_1^+ P_0 + \emptyset = \mathbf{W}_1^+ P_0$, that

$$\begin{aligned} s(N_1) &= s(\mathbf{W}_1^+ P_0) \\ &\leq s(P_0)^d \\ &\leq 1^d \\ &= 1 \end{aligned}$$

and

$$s(P_1) \leq s(P_0)^d + s(N_1) \leq 2,$$

which implies that

$$\xi_1 = s(P_1)s(N_1) \leq 2. \quad (5.9)$$

Thus, the claim holds for $k = 1$. This starts the induction.

Next, assume Equation (5.8) holds for any layer $1 \leq l < k$. We show that it also holds for $l = k$. Indeed,

$$\begin{aligned} \xi_k &\leq 2\xi_{k-1}^{2n_{k-1}} \\ &\stackrel{IH}{\leq} 2 \left(2^{1+\sum_{j=1}^{k-2} \prod_{i=j}^{k-2} 2n_j} \right)^{2n_{k-1}} \\ &= 2 \cdot 2^{2n_{k-1} + \sum_{j=1}^{k-2} \prod_{i=j}^{k-1} 2n_j} \\ &= 2^{1+\sum_{j=1}^{k-1} \prod_{i=j}^{k-1} 2n_j}. \end{aligned}$$

This finishes the proof. \square

Upper bounds for $s(N_L)$ and $s(P_L)$ are now a simple Corollary:

Corollary 5.1.16. *Consider the setting of Proposition 5.1.7 with initial definitions N_0 and P_0 as in Corollary 5.1.8. Then, if the threshold $t_k \neq -\infty$ for all $k \leq l$, the maximum sizes of N_l and P_l are upper bounded by*

$$s(N_l) \leq 2^{n_{l-1}(1+\sum_{j=1}^{l-2} \prod_{i=j}^{l-2} 2n_j)}$$

and

$$s(P_l) \leq 2^{1+n_{l-1}(1+\sum_{j=1}^{l-2} \prod_{i=j}^{l-2} 2n_j)}.$$

Proof. Follows directly from Lemma 5.1.14 and Proposition 5.1.15 \square

If the network's output-dimension is $n_L = 1$ (for example, in the context of Definition 2.1.5 for binary classification), N_L and P_L are vectors containing a single set of dual points each. Thus, we can identify them as sets. The following corollary bounds their size.

Corollary 5.1.17. *Let $\mathcal{N}: \mathbb{R}^d \rightarrow \mathbb{R}$ be a ReLU network with $L \geq 2$ layers. Then we can write*

$$\mathcal{N} = \mathcal{Q}(P) - \mathcal{Q}(N)$$

where $P, Q \subset \mathcal{D}$ are finite sets of dual points whose size is upper bounded by

$$|P|, |N| \leq 2^{n_{L-1}(1+\sum_{j=1}^{L-2} \prod_{i=j}^{L-2} 2n_j)}.$$

Proof. Since $t_L = -\infty$ (i.e., the last layer is linear), Lemma 5.1.14 implies that

$$|P|, |N| \leq \xi_{L-1}^{n_{L-1}}.$$

The claim then follows from Proposition 5.1.15 since $t_l = 0$ for all $l < L$. \square

In the special case where all layers have the same width, the following holds:

Corollary 5.1.18. *Let $\mathcal{N}: \mathbb{R}^d \rightarrow \mathbb{R}$ be network in the sense of Definition 2.1.1 with $L \geq 2$ layers of width $n_l = w \in \mathbb{N}$ for all $l = 1, \dots, L - 1$. Then we can write*

$$\mathcal{N} = \mathcal{Q}(P) - \mathcal{Q}(N)$$

where $P, N \subseteq \mathfrak{D}$ are vectors of finite sets of dual points whose size is upper bounded by

$$s(P), s(N) \leq 2^{w \frac{1-(2w)^{L-2}}{1-2w}} = 2^{\Theta w((2w)^{L-3})}.$$

Proof. The bound follows from Corollary 5.1.17 if we can show that

$$w \left(1 + \sum_{j=2}^{L-2} \prod_{i=j}^{L-2} 2w \right) = w \frac{1 - (2w)^{L-2}}{1 - 2w}.$$

Indeed,

$$\begin{aligned} w \left(1 + \sum_{j=2}^{L-2} \prod_{i=j}^{L-2} 2w \right) &= w \left(1 + \sum_{j=2}^{L-2} (2w)^{L-2-j+1} \right) \\ &\stackrel{(*)}{=} w \left(1 + \sum_{j=1}^{L-3} (2w)^j \right) \\ &= w \sum_{j=0}^{L-3} (2w)^j \\ &= w \frac{1 - (2w)^{L-2}}{1 - 2w} \\ &= w \Theta((2w)^{L-3}) \end{aligned}$$

The equality (*) follows from inserting the upper and lower summation bounds and comparing terms. The last two equalities identify the geometric series and apply known asymptotic behavior. \square

Remark 5.1.19. The following is an outlook and requires results from later on in this work. The reader may want to either read Chapter 6 and Chapter 7 or just consult the central results in Corollary 6.2.3 and Corollary 7.1.8 before reading this remark.

Using these results, we can derive upper bounds on the complexity realizable by a ReLU network $\mathcal{N} = \mathcal{Q}(P) - \mathcal{Q}(N): \mathbb{R}^d \rightarrow \mathbb{R}$ with the first $L - 1$ layers of width w and the last layer of width 1.

In particular, by Corollary 7.1.8, the number of affine regions defined by \mathcal{N} is upper bounded by the cardinality of the set $\mathcal{U}^*(P \diamond N)$ (this bound is tight for random networks by Conjecture 7.1.10). By Corollary 5.1.18,

$$|\mathcal{U}^*(P \diamond N)| \leq |P \diamond N| \leq |P| \cdot |N| \leq 2^{\Theta((2w)^{L-2})}. \quad (5.10)$$

Comparing this to the bound $\mathcal{O}(d^{w(L-1)})$ for the maximum number of affine regions defined by a ReLU network with L layers of width w (see Section 1.2), we conclude that the bound in Equation (5.10) on the number of affine regions is loose.

Similarly, Corollary 5.1.18 together with Corollary 6.2.3 provides the loose upper bound

$$2 \max(|P|, |N|) \leq 2^{1+\Theta((2w)^{L-3})}$$

on the number of linear pieces in the decision boundary of \mathcal{N} .

To understand why these bounds are so loose, note that only the points in the upper convex hulls of P and N contribute to the CPA functions $\mathcal{Q}(P)$ and $\mathcal{Q}(N)$ (see Corollary 3.3.3).

This leads to the following conjecture:

Conjecture 5.1.20. By Corollary 3.3.3, a ReLU network (P_L, N_L) is already uniquely characterized by the upper convex hulls of P_L and N_L . We conjecture that, usually,

$$|\mathcal{U}^*(P)_i| \ll |P_i| \quad \text{and} \quad |\mathcal{U}^*(N)_i| \ll |N_i| \quad \text{for all } i = 1, \dots, L.$$

This conjecture is, for example, supported by our analysis of Telgarsky’s sawtooth function (Proposition 8.2.6 and Figure 8.1).

5.2 Neural Networks and Tropical Geometry

In Section 4.2, we established a relationship between concepts from affine and tropical geometry. Consequently, the findings from Section 5.1 on relating ReLU networks to DCPA functions can be re-formulated in the setting of tropical geometry. In this section, first translate the main duality result in Proposition 5.1.7 to the tropical setting, with the goal of gaining deeper insights and becoming more familiar with this framework.

Similarly to Chapter 4, the tropical considerations in this chapter are not needed for the remainder of this work as all results (except for Corollary 5.2.5) have already been developed in the affine setting. We nevertheless provide this section for the sake of completeness and to make the connection to tropical geometry more explicit.

Definition 5.2.1 (Tropical Rational Signomial Map). A *tropical rational signomial map* is a function $\mathbb{R}^d \rightarrow \mathbb{R}^p$ for some $p \geq 1$ which is a tropical rational signomial in every coordinate.

The following proposition is the translation of the main duality result. It is a generalization of [12, Proposition 5.1], as it does not require integer-valued weights.

Proposition 5.2.2 (NN as Tropical Signomial Rational Functions). *Assume that a neural network in the sense of Definition 2.1.1 can be written up to layer $l - 1$ as a tropical rational signomial map $\mathbf{a}_l(\mathbf{x}) = F_l(\mathbf{x}) \odot G_l(\mathbf{x})$ for some tropical signomial maps F_l and G_l ¹. Then,*

¹Here, we think of \mathbf{a}_l as a function of the network input \mathbf{x}

after writing $\mathbf{W}_l = \mathbf{W}_l^+ - \mathbf{W}_l^-$ using matrices \mathbf{W}_l^+ and \mathbf{W}_l^- with non-negative entries, also the network up to l 'th layer can be written as a tropical rational signomial map

$$\mathbf{a}_{l+1} = F_{l+1} \otimes G_{l+1}.$$

The tropical signomial maps are given by

$$\begin{aligned} G_{l+1} &= \mathbf{W}_{l+1}^+ G_l + \mathbf{W}_{l+1}^- F_l \\ F_{l+1} &= \max\{\mathbf{W}_{l+1}^+ F_l + \mathbf{W}_{l+1}^- G_l + b, G_{l+1} + t\}. \end{aligned}$$

Writing $f_i^{(l)}$ and $g_i^{(l)}$ for the i th coordinate of F_l and G_l , the recurrence takes the form

$$\begin{aligned} g_i^{(l+1)} &= \left[\bigodot_{j=1}^{n_l} \left(f_j^{(l)} \odot w_{ij}^- \right) \right] \odot \left[\left(g_j^{(l)} \right) \odot w_{ij}^+ \right] \\ f_i^{(l+1)} &= \left\{ \left[\bigodot_{j=1}^{n_l} \left(f_j^{(l)} \right) \odot w_{ij}^+ \right] \odot \left[\left(g_j^{(l)} \right) \odot w_{ij}^- \right] \odot b_i \right\} \oplus \left(g_i^{(l+1)} \odot t_i \right) \end{aligned}$$

where we write $(\mathbf{W}_{l+1}^+)_{ij} = w_{ij}^+$ and $(\mathbf{W}_{l+1}^-)_{ij} = w_{ij}^-$.

Proof. A calculation confirms that the above recursion corresponds to the recursion in Proposition 5.1.7 after identifying

$$\begin{aligned} F_l &= \mathcal{Q}(P_{l-1}) \\ G_l &= \mathcal{Q}(N_{l-1}). \end{aligned}$$

□

The following corollary is the analogue to Corollary 5.1.8, which takes place in the affine setting:

Corollary 5.2.3 (Tropical Characterization of Neural Networks). *Every network \mathcal{N} in the sense of Definition 2.1.1 can be written as a tropical rational signomial map. That is,*

$$\mathcal{N} = F \otimes G$$

where F and G are tropical signomial maps.

One can show an even stronger result, namely that we have already found all binary classification networks in the sense of Definition 2.1.5 as tropical rational signomials $\mathbb{R}^d \rightarrow \mathbb{R}$ with finite coefficients. The following theorem is a generalization of [12, Theorem 5.4.i)], since it does not assume integer weights.

Theorem 5.2.4 (Tropical Equivalence). *Let $\mathcal{N} : \mathbb{R}^d \rightarrow \mathbb{R}$ be a function. Then \mathcal{N} is a tropical rational signomial with finite coefficients if and only if it is a neural network in the sense of Definition 2.1.5.*

Proof. The “if” follows from Corollary 5.2.3. We are left to show the “only if”. In this proof we identify finite coefficient tropical signomials with CPA functions in the sense of Proposition 4.2.2.

We first claim that every finite-coefficient tropical signomial $\bigoplus_{i=1}^n b_i \odot x^{\odot\alpha_i} \in \mathbb{T}_{\mathbb{R}}^{\text{fin}}[x_1, \dots, x_d]$ with n tropical summands can be written as a ReLU network with n layers and a final linear layer, i.e.,

$$\bigoplus_{i=1}^n b_i \odot x^{\odot\alpha_i}(\mathbf{x}) = \rho_{-\infty} \circ f_n \circ \rho_0 \circ \dots \circ \rho_0 \circ f_1(\mathbf{x}) \quad (5.11)$$

with affine functions $f_l: \mathbb{R}^{n_{l-1}} \rightarrow \mathbb{R}^{n_l}$.

We show Equation (5.11) by induction on n . The base-case follows readily:

$$\begin{aligned} b_1 x^{\odot\alpha_1}(\mathbf{x}) &= \langle \alpha_1, \mathbf{x} \rangle + b_1 \\ &= \rho_{-\infty}(\langle \alpha_1, \mathbf{x} \rangle + b_1) \\ &= \rho_{-\infty} \circ f_1(\mathbf{x}) \end{aligned}$$

where $f_1 = f_{\alpha_1, b}$. This establishes the base-case.

Assume Equation (5.11) holds for all finite coefficient tropical signomials with less than n tropical summands. We now show that it also holds for n tropical summands.

Let $\bigoplus_{i=1}^n b_i \odot x^{\odot\alpha_i} \in \mathbb{T}_{\mathbb{R}}^{\text{fin}}[x_1, \dots, x_d]$ be a finite coefficient tropical signomial with n tropical summands.

By the induction hypothesis we can find two ReLU network representations

$$\begin{aligned} p(\mathbf{x}) &= \bigoplus_{i=1}^{n-1} b_i \odot x^{\odot\alpha_i}(\mathbf{x}) \\ q(\mathbf{x}) &= b_n \odot x^{\odot\alpha_n}(\mathbf{x}). \end{aligned}$$

Define $y(\mathbf{x}) := (p(\mathbf{x}), q(\mathbf{x}))$. Then y can also be expressed by an $(n - 1)$ -layer ReLU network by extending $q(\mathbf{x})$ using linear layers and isolating the subnetworks using zero-weights:

$$y(\mathbf{x}) = \rho_{-\infty} \circ h_{n-1} \circ \rho_0 \circ \dots \circ \rho_0 \circ h_1(\mathbf{x})$$

for affine functions h_i .

Next, note that

$$\begin{aligned} \bigoplus_{i=1}^n b_i \odot x^{\odot\alpha_i}(\mathbf{x}) &= \max\{p(\mathbf{x}), q(\mathbf{x})\} \\ &= \max\{p(\mathbf{x}) - q(\mathbf{x}), 0\} + q(\mathbf{x}) \\ &= \rho_{-\infty}[\rho_0(p(\mathbf{x}) - q(\mathbf{x})) + \rho_0(q(\mathbf{x})) - \rho_0(-q(\mathbf{x}))] \\ &= \rho_{-\infty} \circ e_n \circ \rho_0 \circ g_n(y(\mathbf{x})) \end{aligned}$$

where e_n is the linear function

$$e_n(\mathbf{x}) = \begin{pmatrix} 1 \\ 1 \\ -1 \end{pmatrix} \mathbf{x}$$

and g_n is the linear function

$$g_n(\mathbf{x}) = \begin{pmatrix} 1 & -1 \\ 0 & 1 \\ 0 & -1 \end{pmatrix} \mathbf{x}.$$

We obtain Equation (5.11) by defining

$$f_i = \begin{cases} h_i & i = 1, \dots, n-2 \\ g_n \circ h_{i-1} & i = n-1 \\ e_n & i = n \end{cases}$$

It remains to show that we can also express a signomial tropical quotient $\varphi \oslash \psi$ as a ReLU network where φ and ψ are tropical signomials with n and m tropical summands respectively. By Equation (5.11), both φ and ψ can be represented as ReLU networks. Again, fuse these two representation into a single ReLU network representation of $z(\mathbf{x}) = (\varphi(\mathbf{x}), \psi(\mathbf{x}))$.

Next, using the fact that $x = x^+ - x^-$, $\varphi \oslash \psi$ can be written

$$\begin{aligned} (\varphi \oslash \psi)(\mathbf{x}) &= \rho_{-\infty}(\rho_0(\varphi(\mathbf{x})) - \rho_0(-\varphi(\mathbf{x})) + \rho_0(-\psi(\mathbf{x})) - \rho_0(\psi(\mathbf{x}))) \\ &= \rho_{-\infty} \circ j_n \circ \rho_0 \circ k_n(z(\mathbf{x})) \end{aligned}$$

where j_n is the linear function

$$j_n(\mathbf{x}) = \begin{pmatrix} 1 \\ -1 \\ 1 \\ -1 \end{pmatrix} \mathbf{x}$$

and k_n is the linear function

$$k_n(\mathbf{x}) = \begin{pmatrix} 1 & 0 \\ -1 & 0 \\ 0 & -1 \\ 0 & 1 \end{pmatrix} \mathbf{x}$$

Analogously to the proof of Equation (5.11), this gives a ReLU representation of $\varphi \oslash \psi$. □

Translated to the setting of affine geometry, the result of Theorem 5.2.4 result reads:

Corollary 5.2.5. *Let $\mathcal{N}: \mathbb{R}^d \rightarrow \mathbb{R}$ be a map. Then \mathcal{N} is a DCPA function if and only if it is a ReLU binary classification network in the sense of Definition 2.1.1.*

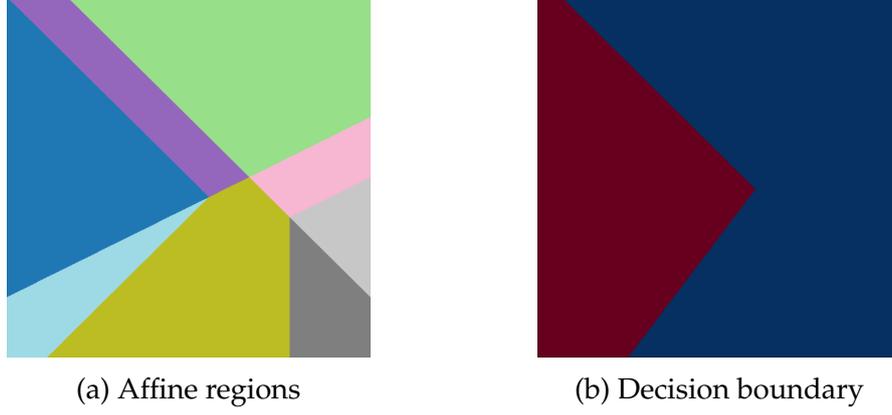


Figure 5.1: Figure (a) shows the affine regions defined by the network defined in Equation (5.12). Figure (b) shows its decision boundary. Negatively classified regions (threshold at 0) are colored red and positively classified regions are colored blue.

Example 5.2.6. In this example, we construct the dual representation of a toy example in two dimensions. Throughout this work, we will revisit and use this example to explain various aspect of the duality result.

Specifically, consider the 3 layer network

$$\mathcal{N}: \mathbb{R}^2 \rightarrow \mathbb{R} \quad (5.12)$$

$$\mathcal{N}(x) = \mathbf{W}_3 \rho_0 (\mathbf{W}_2 \rho_0 (\mathbf{W}_1 \mathbf{x} + \mathbf{b}_1) + \mathbf{b}_2) + b_3 \quad (5.13)$$

where

$$\mathbf{W}_1 = \begin{pmatrix} -1 & -1 \\ 1 & -2 \end{pmatrix}, \mathbf{b}_1 = \begin{pmatrix} 1 \\ -1 \end{pmatrix}, \mathbf{W}_2 = \begin{pmatrix} -1 & 2 \\ 2 & -1 \end{pmatrix}, \mathbf{b}_2 = \begin{pmatrix} 1 \\ 2 \end{pmatrix}, \mathbf{W}_3 = (3, -1), b_3 = 2.$$

We use Proposition 5.1.7 to iteratively construct the dual representation of \mathcal{N} , starting with $P_0 = (\{(1, 0, 0)\}, \{(0, 1, 0)\})$ and $N_0 = (\emptyset)$, as in Corollary 5.1.8.

After the first layer, the dual representation of N_1 can be computed as

$$N_1 = \mathbf{W}_1^- P_0 \diamond \mathbf{W}_1^+ N_0 = \begin{pmatrix} 1 & 1 \\ 0 & 2 \end{pmatrix} P_0$$

and thus

$$(N_1)_1 = \bigoplus_{j=1}^2 (\mathbf{W}_1^-)_{1j} (P_0)_j = 1 \{((1, 0, 0))\} \diamond 1 \{(0, 1, 0)\} = \{(1, 1, 0)\}$$

$$(N_1)_2 = \bigoplus_{j=1}^2 (\mathbf{W}_1^-)_{2j} (P_0)_j = 0 \{((1, 0, 0))\} \diamond 2 \{(0, 1, 0)\} = \{(0, 2, 0)\},$$

which implies

$$N_1 = (\{(1, 1, 0)\}, \{(0, 2, 0)\}).$$

Similarly,

$$P_1 = \mathbf{W}_1^+ P_0 \diamond \mathbf{W}_1^- N_0 \boxplus \mathbf{b}_1 \cup N_1 = \begin{pmatrix} 0 & 0 \\ 1 & 0 \end{pmatrix} P_0 \boxplus \mathbf{b}_1 \cup N_1$$

and thus

$$\begin{aligned} (P_1)_1 &= \bigoplus_{j=1}^2 (\mathbf{W}_1^+)_{1j} (P_0)_j \boxplus (\mathbf{b}_1)_1 \cup (N_1)_1 = 0 \{((1, 0, 0))\} \diamond 0 \{(0, 1, 0)\} \boxplus 1 \cup (N_1)_1 \\ &= \{(0, 0, 1), (1, 1, 0)\} \\ (P_1)_2 &= \bigoplus_{j=1}^2 (\mathbf{W}_1^+)_{2j} (P_0)_j \boxplus (\mathbf{b}_1)_2 \cup (N_1)_2 = 1 \{((1, 0, -1))\} \diamond 0 \{(0, 1, 0)\} \boxplus -1 \cup (N_1)_2 \\ &= \{(1, 0, 0), (0, 2, 0)\}, \end{aligned}$$

which implies

$$P_1 = (\{(0, 0, 1), (1, 1, 0)\}, \{(0, 2, 0), (1, 0, -1)\}).$$

After repeating these steps for layer 2 and 3 (with a slight adaptation for the last linear layer as in Proposition 5.1.7), one arrives at the following final dual representation of $\mathcal{N} = \mathcal{Q}(P) - \mathcal{Q}(N)$:

$$N = \{(3, 17, 4), (2, 16, 5), (5, 19, 2), (3, 14, 2), (2, 16, 3), (5, 19, 0), (6, 17, -1), (0, 14, 7)\} \quad (5.14)$$

$$P = \{(2, 16, 5), (5, 19, 2), (5, 19, 5), (11, 7, -1), (12, 5, -2), (3, 14, 4), (6, 17, 1), (6, 17, 4)\} \quad (5.15)$$

Here, $N := N_3$ and $P := P_3$. Additionally, note that we have identified the one-dimensional vectors of sets N_3 and P_3 with their only entry.

By Corollary 3.3.3, the CPA functions $\mathcal{Q}(N)$ and $\mathcal{Q}(P)$ are uniquely identified by the upper convex hulls of N and P , i.e.

$$\mathcal{Q}(N) = \mathcal{Q}(\mathcal{U}^*(N)), \quad \mathcal{Q}(P) = \mathcal{Q}(\mathcal{U}^*(P)).$$

This allows restricting our attention to subsets of N and P . Specifically, the upper convex hull points can be determined² as

$$\begin{aligned} \mathcal{U}^*(N) &= \{(5, 19, 2), (3, 14, 2), (6, 17, -1), (0, 14, 7)\} \\ \mathcal{U}^*(P) &= \{(2, 16, 5), (3, 14, 4), (5, 19, 5), (12, 5, -2)\}. \end{aligned}$$

Figure 5.2 contains a plot of the dual representation of this toy example, as well as the upper convex hulls.

²We use SciPy to do so, see the code for more details.

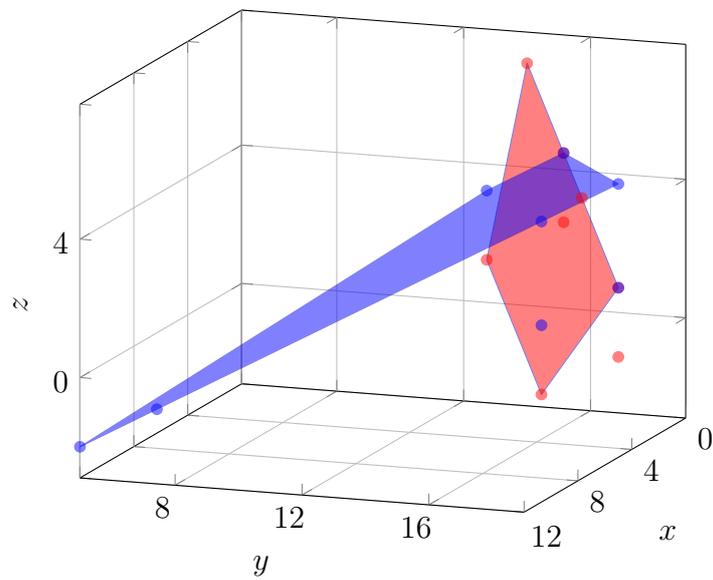


Figure 5.2: Two-dimensional toy-example defined in Equation (5.12). Red points correspond to N , blue points are P . The red polygon is $\mathcal{U}(N)$, the blue polygon is $\mathcal{U}(P)$. Note that, in theory, both $\mathcal{U}(N)$ and $\mathcal{U}(P)$ are polyhedral complexes, i.e., they can consist of multiple facets.

Chapter 6

Characterizing the Decision Boundary

We saw in Chapter 3 how to identify neural networks as DCPA functions using affine geometry. In Chapter 4, we translated this result to tropical geometry, where it reads that every neural network can be written as a tropical rational signomial map.

In this chapter, we use that result to characterize the decision boundary of ReLU binary classification networks. This will eventually allow counting the linear pieces inside the decision boundary.

Throughout this chapter, let $S \subseteq \mathcal{D}$ be a set of dual points whose upper convex hull has vertices

$$\mathcal{U}^*(S) = \{s_1, \dots, s_n\} = \{(\mathbf{a}_1, b_1), \dots, (\mathbf{a}_n, b_n)\}.$$

Furthermore, given a set of indices $I \subseteq \{1, \dots, n\}$, we introduce the short-hand notation

$$S_I := \{s_i \mid i \in I\}$$

for the subset of S indexed by I .

6.1 Identifying Cells and Faces

Towards quantifying the decision boundary complexity, we will establish in Theorem 6.1.7 a bijection between the k -cells in $\mathcal{T}(S)$ and the $(d - k)$ -faces in $\mathcal{U}(S)$. In Section 6.2, we will use this result to translate the decision boundary to dual space, where its complexity is easier to quantify.

Let $\sigma \in \mathcal{T}(S)$ be a cell in the tessellation induced by S . By Definition 3.4.1 and Corollary 3.3.3, σ is the solution of a system of linear inequalities and equalities:

$$\begin{cases} \mathcal{R}(s_i)(\mathbf{x}) = \mathcal{R}(s_{i'}) (\mathbf{x}) \quad \forall i, i' \in I_{=}^{\sigma} \\ \mathcal{R}(s_i)(\mathbf{x}) \geq \mathcal{R}(s_j)(\mathbf{x}) \quad \forall i \in I_{=}^{\sigma}, j \in I_{+}^{\sigma}, \end{cases} \quad (6.1)$$

where $I_{=}^{\sigma}$ and I_{+}^{σ} form a disjoint partition of $\{1, 2, \dots, n\}$. W.l.o.g, this partition can be chosen in such a way that no index can be moved from I_{+}^{σ} to $I_{=}^{\sigma}$ without altering the solution space.

Using the definition that $\mathcal{R}(s_i)(\mathbf{x}) = \langle \mathbf{a}_i, \mathbf{x} \rangle + b_i$ for all $s_i \in S$, System (6.1) can be re-written as a system of linear inequalities:

Definition 6.1.1 (Cells as System of Linear Inequalities). Any cell $\sigma \in \mathcal{T}(S)$ can be written as a system of linear inequalities and equalities $\sigma = \{\mathbf{A}_{\underline{=}}^\sigma \mathbf{x} = \mathbf{b}_{\underline{=}}^\sigma\} \cap \{\mathbf{A}_{\underline{+}}^\sigma \mathbf{x} \geq \mathbf{b}_{\underline{+}}^\sigma\}$ in the following way.

Fix a dual point $s_{k_\sigma} \in S_{I_\sigma^\sigma}$. The matrix $\mathbf{A}_{\underline{=}}^\sigma \in \mathbb{R}^{|I_\sigma^\sigma|, d}$ containing the equality constraints has as its rows the vectors $(\mathbf{a}_{k_\sigma} - \mathbf{a}_i | i \in I_\sigma^\sigma)$ and the corresponding vector $\mathbf{b}_{\underline{=}}^\sigma \in \mathbb{R}^{|I_\sigma^\sigma|}$ has entries $(b_i - b_{k_\sigma} | i \in I_\sigma^\sigma)$ (in the same order).

Similarly, the matrix $\mathbf{A}_{\underline{+}}^\sigma \in \mathbb{R}^{|I_\sigma^\sigma|, d}$ containing the inequality constraints has as its rows the vectors $(\mathbf{a}_{k_\sigma} - \mathbf{a}_i | i \in I_\sigma^\sigma)$ and the corresponding vector $\mathbf{b}_{\underline{+}}^\sigma \in \mathbb{R}^{|I_\sigma^\sigma|}$ has entries $(b_i - b_{k_\sigma} | i \in I_\sigma^\sigma)$.

Remark 6.1.2. The joint system of linear equalities and linear inequalities in Definition 6.1.1 can be translated to a system of just inequalities

$$\sigma = \{\mathbf{A}^\sigma \mathbf{x} \geq \mathbf{b}^\sigma\} = \{\mathbf{A}_{\underline{=}}^\sigma \mathbf{x} = \mathbf{b}_{\underline{=}}^\sigma\} \cap \{\mathbf{A}_{\underline{+}}^\sigma \mathbf{x} \geq \mathbf{b}_{\underline{+}}^\sigma\}$$

by re-writing every equality as two inequalities. The resulting matrix $\mathbf{A}^\sigma \in \mathbb{R}^{2|I_\sigma^\sigma|, d}$ contains the rows of $\mathbf{A}_{\underline{=}}^\sigma$, as well as their negatives, and the rows of $\mathbf{A}_{\underline{+}}^\sigma$. The vector $\mathbf{b}^\sigma \in \mathbb{R}^{2|I_\sigma^\sigma|}$ can be constructed analogously.

Using this representation of σ as a system of linear inequalities, the following proposition describes the dimension of σ :

Proposition 6.1.3. *Let $S \subseteq \mathcal{D}$ be a finite set of dual points and $\sigma \in \mathcal{T}(S)$ be a cell in the tessellation induced by S . Then*

$$\dim \sigma = d - \text{rank } \mathbf{A}_{\underline{=}}^\sigma. \quad (6.2)$$

Proof. This follows directly from Lemma 2.3.10 and since $\mathbf{A}_{\underline{=}}^\sigma$ contains all of the implicit equality constraints by construction of I_σ^σ and I_σ^σ . \square

As a next step, we aim to understand how the cell σ looks in dual space. We start with the following proposition:

Proposition 6.1.4. *Let $\sigma \in \mathcal{T}(S)$ be a cell. Then the convex hull $\mathcal{C}(S_{I_\sigma^\sigma})$ is a face in $\mathcal{U}(S)$ of dimension*

$$\dim \mathcal{C}(S_{I_\sigma^\sigma}) = \text{rank } \mathbf{A}_{\underline{=}}^\sigma. \quad (6.3)$$

Proof. We first show that $\mathcal{C}(S_{I_\sigma^\sigma})$ is contained in a face of dimension $\text{rank } \mathbf{A}_{\underline{=}}^\sigma$. As a second step, we show that it is actually equal to such a face.

For simplicity, abbreviate $\zeta := \mathcal{C}(S_{I_\sigma^\sigma})$. By construction, we know that $S_{I_\sigma^\sigma} \subseteq \mathcal{U}^*(S)$. It is thus left to show that

- i) ζ is a face in $\mathcal{U}(S)$
- ii) $\dim \zeta = \text{rank } \mathbf{A}_{\underline{=}}^\sigma$.

We start with *i*). Assume towards a contradiction there existed a set of positive scalars $\{\alpha_1, \dots, \alpha_{|S_{I_{\underline{\sigma}}}|}\}$ satisfying $\sum_{j=1}^{|S_{I_{\underline{\sigma}}}|} \alpha_j = 1$ such that $\sum_{j \in I_{\underline{\sigma}}} \alpha_j \mathcal{R}(s_j) \notin |\mathcal{U}(S)|$ (the support of the polyhedral complex $\mathcal{U}(S)$, see Definition 2.3.5).

Then, by Equation (3.2),

$$\sum_{j \in I_{\underline{\sigma}}} \alpha_j \mathcal{R}(s_j) < \mathcal{Q}(S)$$

But, for any $\mathbf{x} \in \sigma$,

$$\begin{aligned} \mathcal{Q}(S)(\mathbf{x}) &\geq \mathcal{R}(s_i)(\mathbf{x}) \\ &= \sum_{j \in I_{\underline{\sigma}}} \alpha_j \mathcal{R}(s_j)(\mathbf{x}) \\ &= \sum_{j \in I_{\underline{\sigma}}} \alpha_j \mathcal{R}(s_j)(\mathbf{x}) \\ &\stackrel{3.2.2}{=} \mathcal{R}\left(\sum_{j \in I_{\underline{\sigma}}} \alpha_j s_j\right)(\mathbf{x}). \end{aligned}$$

This is a contradiction. We conclude *i*).

We now show *ii*). It holds that

$$\begin{aligned} \dim \zeta &= \dim \mathcal{C}(S_{I_{\underline{\sigma}}}) \\ &\stackrel{2.2.7}{=} \dim \left(\text{span} \left(\{s_{k_\sigma} - s_j \mid j \in I_{\underline{\sigma}} \setminus \{k_\sigma\}\} \right) \right) \\ &\stackrel{(*)}{=} \text{rank}[\mathbf{A}_{\underline{\sigma}}^\sigma - \mathbf{b}_{\underline{\sigma}}^\sigma] \\ &\stackrel{(**)}{=} \text{rank } \mathbf{A}_{\underline{\sigma}}^\sigma. \end{aligned}$$

Indeed, equality $(*)$ follows by identifying $[\mathbf{A}_{\underline{\sigma}}^\sigma - \mathbf{b}_{\underline{\sigma}}^\sigma]$ with the matrix obtained by stacking the points in $s_{k_\sigma} - S_{I_{\underline{\sigma}}}$ on top of each other. To see why equality $(**)$ holds, note that the system $\mathbf{A}_{\underline{\sigma}}^\sigma \mathbf{x} = \mathbf{b}_{\underline{\sigma}}^\sigma$ has a solution since σ is non-empty. But then also the system $\mathbf{A}_{\underline{\sigma}}^\sigma \mathbf{x} = -\mathbf{b}_{\underline{\sigma}}^\sigma$. Hence, $(**)$ follows.

This shows *ii*) and concludes the first step.

For the second step, we show that ζ is already a face in $\mathcal{U}(S)$, not just contained in one. Assume towards a contradiction it was not and let ξ be the smallest face containing ζ . Then there would exist an $s_j \in \xi \cap \mathcal{U}^*(S)$ with $j \in I_{\underline{\sigma}}^+$ such that

$$\dim \zeta = \dim \mathcal{C}(S_{I_{\underline{\sigma}}} \cup \{s_j\}).$$

By Lemma 2.2.7, this would imply that

$$\dim \zeta = \dim \mathcal{C}(S_{I_{\underline{\sigma}}} \cup s_j) = \dim \left(\text{span} \left(\{s_{k_\sigma} - s_i \mid i \in I_{\underline{\sigma}} \setminus \{k_\sigma\}\} \cup \{s_{k_\sigma} - s_j\} \right) \right)$$

We claim that $\{s_j - s_{k_\sigma}\}$ is linearly independent of $\{s_i - s_{k_\sigma} \mid i \in I_{\underline{\sigma}} \setminus \{k_\sigma\}\}$, which implies

$$\dim \zeta > \dim \left(\text{span} \left(\{s_{k_\sigma} - s_i \mid i \in I_{\underline{\sigma}} \setminus \{k_\sigma\}\} \right) \right) = \text{rank } \mathbf{A}_{\underline{\sigma}}^\sigma,$$

a contradiction to *ii*).

Indeed, assume towards a contradiction that linear independence did not hold. Then there would exist a set $\{\alpha_i\} \subseteq \mathbb{R}$ such that

$$s_{k_\sigma} - s_j = \sum_i \alpha_i (s_{k_\sigma} - s_i)$$

and thus, in particular, by linearity of \mathcal{R} ,

$$(\mathcal{R}(s_{k_\sigma}) - \mathcal{R}(s_j))(\mathbf{x}) = \sum_i \alpha_i (\mathcal{R}(s_{k_\sigma}) - \mathcal{R}(s_i))(\mathbf{x}) = 0 \quad \forall \mathbf{x} \in \sigma.$$

This would imply that s_j is an equality constraint for σ , a contradiction. \square

The last proposition tells us that $\mathcal{C}(S_{\underline{\sigma}})$ is a face in $\mathcal{U}(S)$. The next proposition uses the face $\mathcal{C}(S_{I_{\underline{\sigma}}})$ to explain how σ translates to dual space. But first, we define what it means for an affine function to be tangent to an upper convex hull:

Definition 6.1.5. Given an affine function $f: \mathbb{R}^d \rightarrow \mathbb{R}$, we say that f is *tangent* to the upper convex hull $\mathcal{U}(S)$ if

- i) f lies above $\mathcal{U}(S)$, i.e., $f \succeq \mathcal{U}(S)$
- ii) and the graph of f intersects the upper convex hull, i.e., $\text{graph}(f) \cap \mathcal{U}(S) \neq \emptyset$.

In this case, we write $f \parallel \mathcal{U}(S)$.

Proposition 6.1.6 ([18, Proposition 18]). *Let $\sigma \in \mathcal{T}(S)$ be a cell in the tessellation induced by S . Then there is a one-to-one-correspondence between points in σ and dual planes tangent to the upper convex hull of S which contain the face $\mathcal{C}(S_{I_{\underline{\sigma}}})$:*

$$\{\mathbf{x} \in \sigma\} \xleftrightarrow{1:1} \{f \in \text{Aff}_{\mathbb{D}}(\mathbb{d}) \mid f \parallel \mathcal{U}(S) \text{ and } f \supseteq \mathcal{C}(S_{I_{\underline{\sigma}}})\}. \quad (6.4)$$

Similarly, every face in $\mathcal{U}(S)$ defines a cell in this way.

Proof. By definition, σ is the region in \mathbb{R}^d defined by a system

$$\begin{cases} \mathcal{R}(s_i)(\mathbf{x}) = \mathcal{R}(s_{i'}) (\mathbf{x}) \quad \forall i, i' \in I_{\underline{\sigma}} \\ \mathcal{R}(s_i)(\mathbf{x}) \geq \mathcal{R}(s_j)(\mathbf{x}) \quad \forall i \in I_{\underline{\sigma}}, j \in I_{\underline{\sigma}}^c. \end{cases} \quad (6.5)$$

Using the notation introduced in Definition 2.2.1, this system can be re-written as

$$\begin{cases} (\mathbf{x}, y) \in \mathcal{R}(s_i) \quad \forall i \in I_{\underline{\sigma}} \\ (\mathbf{x}, y) \succeq \mathcal{R}(s_j) \quad \forall j \in I_{\underline{\sigma}}^c \end{cases}$$

(where $y = \mathcal{Q}(S)(\mathbf{x})$). By Proposition 3.2.4, it can be translated to dual space:

$$\begin{cases} \tilde{\mathcal{R}}^{-1}((\mathbf{x}, y)) \ni s_i \quad \forall i \in I_{\underline{\sigma}} \\ \tilde{\mathcal{R}}^{-1}((\mathbf{x}, y)) \succeq s_j \quad \forall j \in I_{\underline{\sigma}}^c. \end{cases} \quad (6.6)$$

The first row in System (6.6) describes dual planes containing $\zeta := \mathcal{C}(S_{I_\sigma})$. We claim that the second row of System (6.6) restricts those to dual planes upper convex hull. Indeed, for such a dual plane to be tangent to the upper convex hull, the following needs to be true:

- i) $\tilde{\mathcal{R}}^{-1}((\mathbf{x}, y)) \succeq \mathcal{U}(S)$
- ii) $\tilde{\mathcal{R}}^{-1}((\mathbf{x}, y)) \cap \mathcal{U}(S) \neq \emptyset$.

But this follows immediately from System (6.6). This shows one implication of the proposition by establishing the one-to-one correspondence.

For the other implication, start with a face $\zeta \in \mathcal{U}(S)$ and let $I \subseteq \{1, 2, \dots, n\}$ be the set indexing the points in S that lie on ζ . Reversing the argument by transitioning from System (6.6) back to System (6.5) yields a cell $\sigma \in \mathcal{T}(S)$ with $I_\sigma = I$. \square

By Proposition 6.1.4 and Proposition 6.1.6, cells in the tessellation induced by S are closely related to faces in the upper convex hull of S . The following proposition makes this relationship precise.

Theorem 6.1.7. *There exists a one-to-one-correspondence between k -cells in $\mathcal{T}(S)$ and $(d-k)$ -faces in $\mathcal{U}(S)$. Specifically, the following map is a bijection:*

$$\Phi: \mathcal{T}_k(S) \xrightarrow{\sim} \mathcal{U}_{d-k}(S) \tag{6.7}$$

$$\sigma \mapsto \mathcal{C}(S_{I_\sigma}). \tag{6.8}$$

Proof. We first show that the map is well-defined. Let $\sigma \in \mathcal{T}_k(S)$ be a k -cell and write $\zeta := \mathcal{C}(S_{I_\sigma})$. By Proposition 6.1.4, ζ is a face in $\mathcal{U}(S)$ of dimension

$$\dim \zeta = \text{rank } \mathbf{A}_\sigma. \tag{6.9}$$

By Proposition 6.1.3,

$$k = \dim \sigma = d - \text{rank } \mathbf{A}_\sigma. \tag{6.10}$$

Thus, $\zeta \in \mathcal{U}_{d-k}(S)$. This shows that Φ is well-defined.

Injectivity of Φ follows from the injectivity of the assignment $\sigma \mapsto I_\sigma$; every cell in $\mathcal{T}(S)$ has a unique set of maximizers.

Finally, surjectivity of Φ can be proved as follows. Let $\zeta \in \mathcal{U}_{d-k}(S)$ be a $(d-k)$ -face.

Clearly, the set $\{f \in \text{Aff}_{\mathfrak{R}}(d) \mid f \parallel \mathcal{U}(S) \text{ and } f \supseteq \zeta\}$ is non-empty. By Proposition 6.1.6, ζ corresponds to a cell $\sigma \in \mathcal{T}(S)$ with $\mathcal{C}(I_\sigma) = \zeta$. By an argument analogous to the one above, $\dim \sigma = k$. This cell is a pre-image of ζ under Φ . \square

6.2 Application to the Decision Boundary

In this section, we use the bijection from Theorem 6.1.7 to characterize the decision boundary of a ReLU binary classification network $\mathcal{N} = \mathcal{Q}(P) - \mathcal{Q}(N): \mathbb{R}^d \rightarrow \mathbb{R}$.

As a quick reminder, the networks decision boundary is given by

$$\mathcal{B} = (\mathcal{Q}(P) - \mathcal{Q}(N))^{-1}(0).$$

Consequently, we are interested in studying zero-sets of DCPA functions. We start with a special case.

Proposition 6.2.1 (Decision Boundary I [18, Proposition 19]). *Let $F = \mathcal{Q}(P)$ and $G = \mathcal{Q}(N)$ be CPA functions $\mathbb{R}^d \rightarrow \mathbb{R}$ for some finite sets of dual points $P, N \subseteq \mathfrak{D}$. Assume that no point of P lies on $\mathcal{U}(N)$ and vice versa. Let D be the zero-set of $F - G$. Then D is the union of precisely those $(d - 1)$ -cells of $\mathcal{T}(P \cup N)$ which (in the sense of Theorem 6.1.7) correspond to edges (i.e., 1-faces) of $\mathcal{U}(P \cup N)$ with one end in P and the other end in N .*

Proof. For ease of notation, enumerate $N = \{n_1, \dots, n_m\}$ and $P = \{p_1, \dots, p_k\}$ with $m, k \in \mathbb{N}$.

Fix $\mathbf{x} \in \mathbb{R}^d$. Then \mathbf{x} lies in D if and only if $\mathcal{Q}(P)(\mathbf{x}) = \mathcal{Q}(N)(\mathbf{x})$, which by definition means that $\max_{1 \leq i \leq m} \mathcal{R}(n_i)(\mathbf{x}) = \max_{1 \leq j \leq k} \mathcal{R}(p_j)(\mathbf{x})$. Let $\mathcal{R}(n_i)$ and $\mathcal{R}(p_j)$ be maximizers from both sides of the equation, respectively. Then \mathbf{x} is contained in the cell σ defined as the solution of the system

$$\begin{cases} \mathcal{R}(n_i) = \mathcal{R}(p_j) \\ \mathcal{R}(n_i) \geq \mathcal{R}(n_k) \quad \forall k \in \{1, \dots, m\} \setminus \{i\} \\ \mathcal{R}(p_j) \geq \mathcal{R}(p_k) \quad \forall k \in \{1, \dots, k\} \setminus \{j\}. \end{cases} \quad (6.11)$$

By Proposition 6.1.3, this cell has dimension

$$\dim \sigma = d - \text{rank } \mathbf{A}_{\underline{\sigma}} = d - 1. \quad (6.12)$$

By the bijection in Theorem 6.1.7, such a cell corresponds to a dual 1-face (i.e., an edge) of $\mathcal{U}(P \cup N)$ containing the vertex $n_i \in N$ and the vertex $p_j \in P$. It could indeed contain even more points which lie in the affine subspace spanned by p_j and n_i . See Remark 6.2.4 for more details.

Since $\mathcal{P} \cap \mathcal{U}(N) = N \cap \mathcal{U}(P) = \emptyset$, such a face has to have one end in P and one end in N . The correspondence constructed this way is clearly one-to-one. \square

Proposition 6.2.1 handles the special case that $\mathcal{P} \cap \mathcal{U}(N) = N \cap \mathcal{U}(P) = \emptyset$. The following proposition handles the general case.

Proposition 6.2.2 (Decision Boundary II [18, Proposition 20]). *Let $F = \mathcal{Q}(P)$ and $G = \mathcal{R}(N)$ be CPA functions $\mathbb{R}^d \rightarrow \mathbb{R}$ for some finite sets of dual points $P, N \subseteq \mathfrak{D}$. Let D be the zero-set of $F - G$. Then D is the union of precisely those $(d - 1)$ -cells of $\mathcal{T}(P \cup N)$ which (in the sense of Theorem 6.1.7) correspond to edges of $\mathcal{U}(P \cup N)$ containing points from both P and N .*

Proof. We start the proof as the one for Proposition 6.2.1 up to the identification of the edges in the upper convex hull. Since $\mathcal{P} \cap \mathcal{U}(N)$ and $N \cap \mathcal{U}(P)$ need not be empty, we can not conclude that edges in $\mathcal{U}(P \cup N)$ which contain both points from P and N have to start in P and end in N . Thus, we have to require that the edge contains points from both P and N . \square

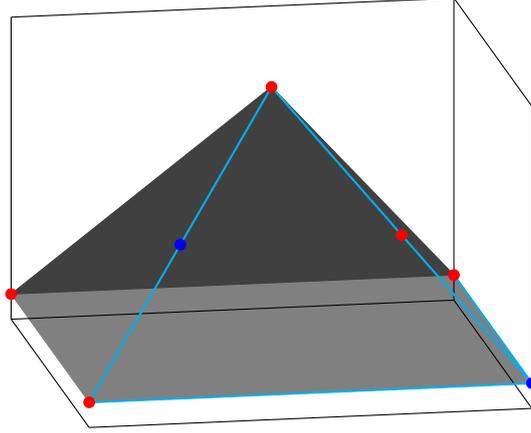


Figure 6.1: An example of points (P, N) (points in N are red, points in P are blue), defining a ReLU network $\mathcal{N} = \mathcal{Q}(P) - \mathcal{Q}(N): \mathbb{R}^2 \rightarrow \mathbb{R}$. There are four edges (light blue) contributing to the decision boundary of \mathcal{N} , since they contain both red and blue points. However, only three of them start and end in different colors.

For completeness, the following corollary applies these findings to neural networks.

Corollary 6.2.3. *Let $\mathcal{Q}(P) - \mathcal{Q}(N): \mathbb{R}^d \rightarrow \mathbb{R}$ be a ReLU binary classification network in the sense of Definition 2.1.5. Then the number of linear pieces in the decision boundary of \mathcal{N} equals the number of edges in $\mathcal{U}(P \cup N)$ containing points from both P and N .*

Remark 6.2.4. The reader may encounter some confusion regarding the proof of Proposition 6.2.1, as well as the necessity for the more general Proposition 6.2.2. We try to clarify matters in this remark, which is accompanied by Figure 6.1.

Fix an $x \in D$ and let $\mathcal{R}(n_i)$ and $\mathcal{R}(p_j)$ be maximizers of $\mathcal{Q}(N)$ and $\mathcal{Q}(P)$, respectively, as in the proof of Proposition 6.2.1. Denote by σ the 1-cell containing x and by ζ the corresponding face.

To clarify the proof of Proposition 6.2.1, specifically Equation (6.12), assume there existed another vertex $n_k \in N$ which is a true convex combination of n_i and p_j , i.e., there exists an $\alpha \in (0, 1)$ such that $n_k = \alpha n_i + (1 - \alpha)p_j$. Then, clearly, n_k lies on ζ .

By linearity of \mathcal{R} , $\mathcal{R}(n_k)$ is also a maximizer of $\mathcal{Q}(N)$ on x . This shows that the system of implicit equalities corresponding to σ is at least made up of n_i, n_k, p_j . Relevant for the dimension of σ , however, is just the rank of $\mathbf{A}_{\sigma}^{\zeta}$, and since n_i, n_k, p_j are affinely dependent, this is not a contradiction.

This should clarify the proof of Proposition 6.2.1. We now want to discuss the difference between Proposition 6.2.1 and Proposition 6.2.2. Generally, there are two different ways to place the three points n_i, n_k, p_j on ζ . In one arrangement, ζ starts in a point from N , ends in a point from P and contains the third point, from N , in its interior. In another setup, ζ starts and ends in N , containing the point from P in its interior (see Figure 6.1 for an example). In both arrangements, one of the three points is a convex combination of the other ones, explaining how all points can contribute a maximizer on x . In both arrangements, the face ζ contains points from both P and N . But Proposition 6.2.1 only allows for the first arrangement, while Proposition 6.2.2 allows for

both.

To end this section, we compare the result of Proposition 6.2.2 to [12, Proposition 6.1.ii)], which is derived in the tropical setting and states that

$$\mathcal{B} \subseteq \mathcal{H}(\mathcal{Q}(N) \oplus \mathcal{Q}(P)),$$

where we identify the $(d - 1)$ -skeleton \mathcal{T}_{d-1} with the tropical hypersurface \mathcal{H} [12, Definition 1], and think of $\mathcal{Q}(N)$ and $\mathcal{Q}(P)$ as tropical signomials

Then, by Lemma 5.1.2.i),

$$\mathcal{H}(\mathcal{Q}(N) \oplus \mathcal{Q}(P)) = \mathcal{H}(\mathcal{Q}(P \cup N)) = \mathcal{T}_{d-1}(P \cup N).$$

Hence, [12, Proposition 6.1.ii)] gives an upper bound on the number of linear pieces in the decision boundary by considering all $(d - 1)$ -cells, not just the ones containing a point from both P and N .

Example 6.2.5. In this subsection, we continue the toy-example $\mathcal{N} = \mathcal{Q}(P) - \mathcal{Q}(N)$ from Example 5.2.6. By Corollary 6.2.3, the number of linear pieces in its decision boundary is the same as the number of edges in $\mathcal{U}(P \cup N)$ containing points from both P and N .

Specifically, one can compute

$$\mathcal{U}^*(P \cup N) = \{(5, 19, 5), (0, 14, 7), (12, 5, -2)\}$$

where $(5, 19, 5), (12, 5, -2) \in \mathcal{U}^*(P)$ and $(0, 14, 7) \in \mathcal{U}^*(N)$ (see Figure 6.2). Thus, there are three edges in $\mathcal{U}(P \cup N)$, two of which contribute to the network's decision boundary since they contain vertices from both P and N . This confirms Figure 5.1b, which shows the decision boundary of \mathcal{N} and confirms that, indeed, it consists of two linear pieces.

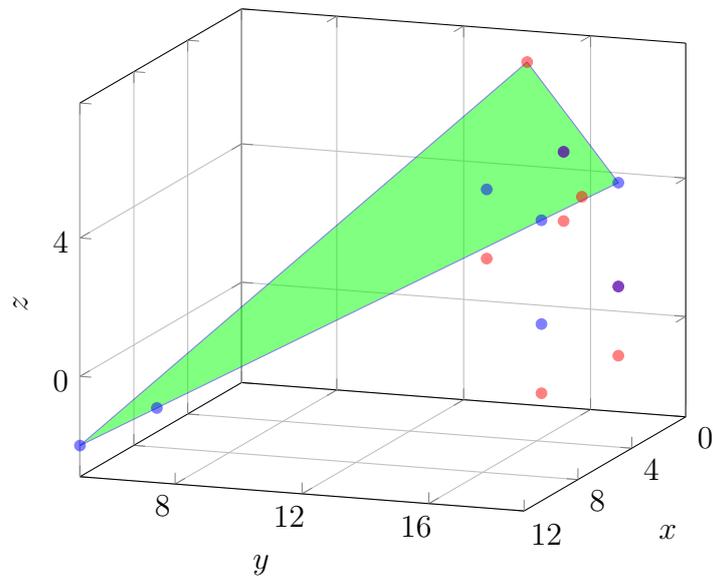


Figure 6.2: Two-dimensional toy-example defined in Equation (5.12). Red points correspond to N , blue points are P . The green polygon is $\mathcal{U}(P \cup N)$. Note that, in theory, $\mathcal{U}(P \cup N)$ and is a polyhedral complex, i.e., it can consist of multiple facets. Note also how there are red and blue points in $\mathcal{U}^*(P \cup N)$, ultimately contributing to the decision boundary.

Chapter 7

Characterizing Affine Regions

In Chapter 6, we used the upper convex hull of $P \cup N$ to characterize the decision boundary of a ReLU binary classification network $\mathcal{Q}(P) - \mathcal{Q}(N)$. In this chapter, we take a similar approach to better understand the k -cells in $\mathcal{T}(P, N)$. In particular, this will enable us to count the network's affine regions.

Any cell $\sigma \in \mathcal{T}(P, N)$ is of the form $\sigma = \sigma' \cap \sigma''$ for some $\sigma' \in \mathcal{T}(P)$ and $\sigma'' \in \mathcal{T}(N)$. In Definition 6.1.1 and Remark 6.1.2, we saw how σ' and σ'' can be expressed as the solution of a system of linear inequalities $\sigma' = \{\mathbf{A}^{\sigma'} \mathbf{x} \geq \mathbf{b}^{\sigma'}\}$, $\sigma'' = \{\mathbf{A}^{\sigma''} \mathbf{x} \geq \mathbf{b}^{\sigma''}\}$. This induces a similar representation for σ :

$$\sigma = \left\{ \begin{bmatrix} \mathbf{A}^{\sigma'} \\ \mathbf{A}^{\sigma''} \end{bmatrix} \mathbf{x} \geq \begin{bmatrix} \mathbf{b}^{\sigma'} \\ \mathbf{b}^{\sigma''} \end{bmatrix} \right\}. \quad (7.1)$$

Analogously to Chapter 6, we now turn our attention to the induced system of implicit equalities.

Definition 7.0.1 (Refined Cells as System of Linear Inequalities). Let $\{\mathbf{A}_{\underline{\underline{I}}}^{\sigma', \sigma''} \mathbf{x} = \mathbf{b}_{\underline{\underline{I}}}^{\sigma', \sigma''}\}$ be the system of implicit inequalities in σ coming from σ' . That is, any row $\mathbf{a}_i^{\sigma'} \in \mathbf{A}_{\underline{\underline{I}}}^{\sigma', \sigma''}$ is also a row in $\mathbf{A}^{\sigma'}$ and satisfies

$$\langle \mathbf{a}_i^{\sigma'}, \mathbf{x} \rangle = \mathbf{b}_i^{\sigma'} \quad \forall \mathbf{x} \in \sigma' \cap \sigma''.$$

We write $I_{\underline{\underline{I}}}^{\sigma', \sigma''}$ for the set indexing these implicit equality constraints. Generally, $I_{\underline{\underline{I}}}^{\sigma'} \subseteq I_{\underline{\underline{I}}}^{\sigma', \sigma''}$, since the latter could contain constraints that only become implicit equality constraints in combination with σ'' (see Figure 7.1 for an example).

Similarly, define the system of implicit equalities in σ coming from σ'' as $\{\mathbf{A}_{\underline{\underline{I}}}^{\sigma'', \sigma'} \mathbf{x} = \mathbf{b}_{\underline{\underline{I}}}^{\sigma'', \sigma'}\}$. That is, any row $\mathbf{a}_i^{\sigma''} \in \mathbf{A}_{\underline{\underline{I}}}^{\sigma'', \sigma'}$ is also a row in $\mathbf{A}^{\sigma''}$ and satisfies

$$\langle \mathbf{a}_i^{\sigma''}, \mathbf{x} \rangle = \mathbf{b}_i^{\sigma''} \quad \forall \mathbf{x} \in \sigma' \cap \sigma''.$$

Again, let $I_{\underline{\underline{I}}}^{\sigma'', \sigma'}$ be the set indexing these implicit equality constraints.

The following proposition describes the dimension of σ in this setup:

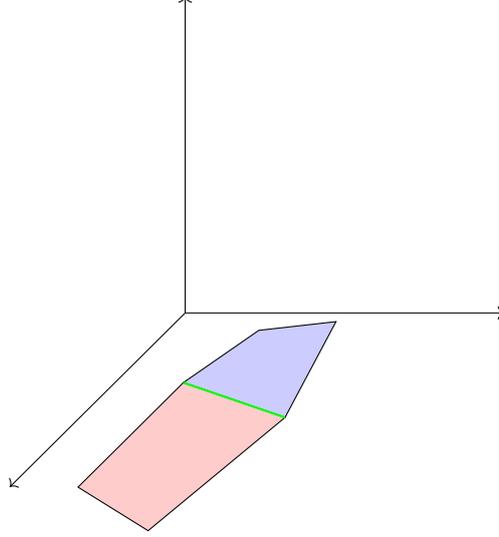


Figure 7.1: Example of a cell σ (green line) in the tessellation $\mathcal{T}(P, N)$, formed by the intersection of a cell $\sigma' \in \mathcal{T}(P)$ (blue) and a cell $\sigma'' \in \mathcal{T}(N)$ (red). Both σ' and σ'' have co-dimension one, each satisfying a single equality constraint. Their intersection imposes an additional equality constraint, resulting in σ having co-dimension 2.

Proposition 7.0.2. *Let $\sigma = \sigma' \cap \sigma'' \in \mathcal{T}(P, N)$ be a cell in the tessellation induced by $\mathcal{Q}(P) - \mathcal{Q}(N)$. Then*

$$\dim \sigma = d - \text{rank} \begin{bmatrix} \mathbf{A}_{=}^{\sigma', \sigma''} \\ \mathbf{A}_{=}^{\sigma'', \sigma'} \end{bmatrix}$$

Proof. By Lemma 2.3.10,

$$\dim \sigma = d - \text{rank} \begin{bmatrix} \mathbf{A}^{\sigma'} \\ \mathbf{A}^{\sigma''} \end{bmatrix}_{=}.$$

The proposition then follows from the fact that

$$\text{rank} \begin{bmatrix} \mathbf{A}^{\sigma'} \\ \mathbf{A}^{\sigma''} \end{bmatrix}_{=} = \text{rank} \begin{bmatrix} \mathbf{A}_{=}^{\sigma', \sigma''} \\ \mathbf{A}_{=}^{\sigma'', \sigma'} \end{bmatrix}.$$

Indeed, both matrices contain the implicit equality constraints for $\sigma = \sigma' \cap \sigma''$. However, the matrix on the right hand side might contain duplicate rows if $I_{=}^{\sigma', \sigma''} \cap I_{=}^{\sigma'', \sigma'} \neq \emptyset$. \square

Like derived for the decision boundary in the previous chapter, the next step is to understand how σ appears in dual space. We begin with the following proposition:

Proposition 7.0.3. *Let $\sigma = \sigma' \cap \sigma'' \in \mathcal{T}_k(P, N)$ be a k -cell in the tessellation induced by $\mathcal{Q}(P) - \mathcal{Q}(N)$. Then*

$$\mathcal{C}(P_{I_{=}^{\sigma', \sigma''}} \diamond N_{I_{=}^{\sigma'', \sigma'}}) \in \mathcal{U}_{d-k}(P \diamond N).$$

Proof. Like in the proof of Proposition 6.1.4, we first show that $\mathcal{C}(P_{I_{=}^{\sigma', \sigma''}} \diamond N_{I_{=}^{\sigma'', \sigma'}})$ is contained in a face in $\mathcal{U}(P \diamond N)$. To do so, however, we take a different approach by first translating σ to dual space. For simplicity, we abbreviate $\zeta := \mathcal{C}(P_{I_{=}^{\sigma', \sigma''}} \diamond N_{I_{=}^{\sigma'', \sigma'}})$.

By definition, σ is the solution of the system

$$\begin{cases} \mathcal{R}(p_i)(\mathbf{x}) = \mathcal{R}(p_{i'}) (\mathbf{x}) & \forall i, i' \in I_{\underline{=}}^{\sigma', \sigma''} \\ \mathcal{R}(p_i)(\mathbf{x}) \succeq \mathcal{R}(p_j)(\mathbf{x}) & \forall i \in I_{\underline{=}}^{\sigma', \sigma''}, j \in I_{+}^{\sigma', \sigma''} \\ \mathcal{R}(n_k)(\mathbf{x}) = \mathcal{R}(n_{k'}) (\mathbf{x}) & \forall k, k' \in I_{\underline{=}}^{\sigma'', \sigma'} \\ \mathcal{R}(n_k)(\mathbf{x}) \succeq \mathcal{R}(n_l)(\mathbf{x}) & \forall k \in I_{\underline{=}}^{\sigma'', \sigma'}, l \in I_{+}^{\sigma'', \sigma'}. \end{cases} \quad (7.2)$$

With the notation introduced in Definition 2.2.1, this system can be re-written as

$$\begin{cases} (\mathbf{x}, y_1) \in \mathcal{R}(p_i) & \forall i \in I_{\underline{=}}^{\sigma', \sigma''} \\ (\mathbf{x}, y_1) \succeq \mathcal{R}(p_j) & \forall j \in I_{+}^{\sigma', \sigma''} \\ (\mathbf{x}, y_2) \in \mathcal{R}(n_k) & \forall k \in I_{\underline{=}}^{\sigma'', \sigma'} \\ (\mathbf{x}, y_2) \succeq \mathcal{R}(n_l) & \forall l \in I_{+}^{\sigma'', \sigma'}, \end{cases}$$

where $y_1 = \mathcal{Q}(P)(\mathbf{x})$ and $y_2 = \mathcal{Q}(N)(\mathbf{x})$. Using Proposition 3.2.4, the system can be translated to dual space:

$$\begin{cases} \check{\mathcal{R}}^{-1}((\mathbf{x}, y_1)) \ni p_i & \forall i \in I_{\underline{=}}^{\sigma', \sigma''} \\ \check{\mathcal{R}}^{-1}((\mathbf{x}, y_1)) \succeq p_j & \forall j \in I_{+}^{\sigma', \sigma''} \\ \check{\mathcal{R}}^{-1}((\mathbf{x}, y_2)) \ni n_k & \forall k \in I_{\underline{=}}^{\sigma'', \sigma'} \\ \check{\mathcal{R}}^{-1}((\mathbf{x}, y_2)) \succeq n_l & \forall l \in I_{+}^{\sigma'', \sigma'}. \end{cases} \quad (7.3)$$

Think of $\check{\mathcal{R}}^{-1}((\mathbf{x}, y_1))$ and $\check{\mathcal{R}}^{-1}((\mathbf{x}, y_2))$ as a pair of parallel hyperplanes, where the former is tangent to $\mathcal{U}(P)$ and contains the cell defined by the first line in System (7.3), and the latter is parallel to $\mathcal{U}(N)$ and contains the cell defined by the third line in System (7.3) (see the proof of Proposition 6.1.6 for details).

View these hyperplanes as subsets of \mathbb{R}^{d+1} . Then their Minkowski sum satisfies

$$\check{\mathcal{R}}^{-1}((\mathbf{x}, y_1 + y_2)) = \check{\mathcal{R}}^{-1}((\mathbf{x}, y_1)) \diamond \check{\mathcal{R}}^{-1}((\mathbf{x}, y_2)) \succeq p + n \quad \forall p \in P, n \in N. \quad (7.4)$$

This means that $\check{\mathcal{R}}^{-1}((\mathbf{x}, y_1 + y_2))$ is tangent to $\mathcal{U}(P \diamond N)$. Indeed, by Definition 6.1.5, we need to check that

- i) $\check{\mathcal{R}}^{-1}((\mathbf{x}, y_1 + y_2))$ lies above $\mathcal{U}(P \diamond N)$, i.e. $\check{\mathcal{R}}^{-1}((\mathbf{x}, y_1 + y_2)) \succeq \mathcal{U}(P \diamond N)$
- ii) $\check{\mathcal{R}}^{-1}((\mathbf{x}, y_1 + y_2)) \cap \mathcal{U}(P \diamond N) \neq \emptyset$.

But *i*) follows directly from Equation (7.4) and *ii*) if we can show that ζ is contained in $\check{\mathcal{R}}^{-1}((\mathbf{x}, y_1 + y_2))$. But this follows from Equation (7.4) and lines 1 and 3 of System 7.3, which imply that

$$\check{\mathcal{R}}^{-1}((\mathbf{x}, y_1 + y_2)) = \check{\mathcal{R}}^{-1}((\mathbf{x}, y_1)) \diamond \check{\mathcal{R}}^{-1}((\mathbf{x}, y_2)) \ni p_i + n_k \quad \forall i \in I_{\underline{=}}^{\sigma', \sigma''}, n_k \in I_{\underline{=}}^{\sigma'', \sigma'}.$$

This shows that ζ is indeed contained inside a face in $\mathcal{U}(S)$ (see Figure 7.2 for a visualization).

Continuing to mimic the proof of Proposition 6.1.4, we next show that

$$\dim \zeta = d - k.$$

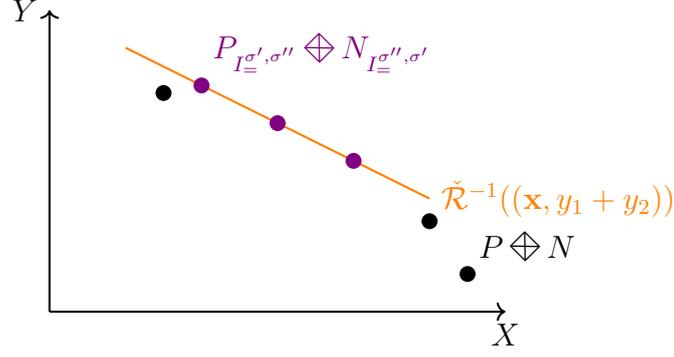


Figure 7.2: Example of a set $P_{I_{\sigma', \sigma''}} \diamond N_{I_{\sigma'', \sigma'}}$ (violet) in $P \diamond N$ (black). If there exists a hyperplane $\check{R}^{-1}((\mathbf{x}, y_1 + y_2))$ tangent to $\mathcal{U}(P \diamond N)$ (orange) containing $\zeta = \mathcal{C}(P_{I_{\sigma', \sigma''}} \diamond N_{I_{\sigma'', \sigma'}})$, then ζ has to be a face in $\mathcal{U}(P \diamond N)$.

But this follows from Lemma 2.2.8, which implies that

$$\begin{aligned}
\dim \zeta &= \dim \mathcal{C}(P_{I_{\sigma', \sigma''}} \diamond N_{I_{\sigma'', \sigma'}}) \\
&= \dim \left(\text{span} \left(\left\{ p_{k_{\sigma'}} - p_i \mid i \in I_{\sigma', \sigma''} \setminus \{p_{k_{\sigma'}}\} \right\} \cup \left\{ n_{k_{\sigma''}} - n_k \mid k \in I_{\sigma'', \sigma'} \setminus \{p_{k_{\sigma''}}\} \right\} \right) \right) \\
&= \text{rank} \begin{bmatrix} \mathbf{A}_{I_{\sigma', \sigma''}}^{\sigma', \sigma''} & -\mathbf{b}_{I_{\sigma', \sigma''}}^{\sigma', \sigma''} \\ \mathbf{A}_{I_{\sigma'', \sigma'}}^{\sigma'', \sigma'} & -\mathbf{b}_{I_{\sigma'', \sigma'}}^{\sigma'', \sigma'} \end{bmatrix} \\
&= \text{rank} \begin{bmatrix} \mathbf{A}_{I_{\sigma', \sigma''}}^{\sigma', \sigma''} \\ \mathbf{A}_{I_{\sigma'', \sigma'}}^{\sigma'', \sigma'} \end{bmatrix} \\
&= d - k,
\end{aligned}$$

where in the last step we used Proposition 7.0.2.

Finally, as in the proof of Proposition 6.1.4, one can show that ζ is already a face in $\mathcal{U}(S)$, by assuming it is not and arriving at a contradiction using Lemma 2.2.8. \square

The last proposition tells us that $\mathcal{C}(P_{I_{\sigma', \sigma''}} \diamond N_{I_{\sigma'', \sigma'}})$ is a face in $\mathcal{U}(P \diamond N)$. As in the previous chapter, the next proposition uses this face to explain how σ looks in dual space:

Proposition 7.0.4 ([18, Proposition 22]). *Let $\sigma = \sigma' \cap \sigma'' \in \mathcal{T}(P, N)$ be a cell in the tessellation induced by $\mathcal{Q}(P) - \mathcal{Q}(N)$. Then there is a one-to-one-correspondence between points in σ and dual planes tangent to the upper convex hull $\mathcal{U}(P \diamond N)$ containing the face $\mathcal{C}(P_{I_{\sigma', \sigma''}} \diamond N_{I_{\sigma'', \sigma'}})$:*

$$\{\mathbf{x} \in \sigma\} \xleftrightarrow{1:1} \{f \in \text{Aff}_{\mathfrak{D}}(d) \mid f \parallel \mathcal{U}(P \diamond N) \text{ and } f \supseteq \mathcal{C}(P_{I_{\sigma', \sigma''}} \diamond N_{I_{\sigma'', \sigma'}})\}. \quad (7.5)$$

Similarly, every face in $\mathcal{U}(P \diamond N)$ corresponds to a cell in this way.

Proof. The construction of the hyperplane $\check{R}^{-1}((\mathbf{x}, y_1 + y_2))$ and the face ζ in the proof of Proposition 7.0.3 is unique for any $\mathbf{x} \in \sigma$, which shows there exists an injection

$$\{\mathbf{x} \in \sigma\} \hookrightarrow \{f \in \text{Aff}_{\mathfrak{D}}(d) \mid f \parallel \mathcal{U}(P \diamond N) \text{ and } f \supseteq \mathcal{C}(P_{I_{\sigma', \sigma''}} \diamond N_{I_{\sigma'', \sigma'}})\}.$$

Next, we show that this is actually a bijection. Let $f_{\mathbf{a},b} \in \text{Aff}_{\mathcal{D}}(d)$ be a hyperplane tangent to $\mathcal{U}(P \diamond N)$ that contains ζ . We need to show that $-\mathbf{a} \in \sigma$ (the additional $-$ comes from the fact that $\tilde{\mathcal{R}}^{-1}((\mathbf{a}, b)) = (-\mathbf{a}, b)$). Assume towards a contradiction that $-\mathbf{a} \notin \sigma$, then there would either exist a $p_i \in P_{I_+^{\sigma', \sigma''}}$ and a $p_k \in P_{I_-^{\sigma', \sigma''}}$ s.t.

$$\mathcal{R}(p_i)(-\mathbf{a}) > \mathcal{R}(p_k)(-\mathbf{a}) \quad (7.6)$$

or an $n_j \in N_{I_+^{\sigma'', \sigma'}}$ and an $n_k \in N_{I_-^{\sigma'', \sigma'}}$ s.t.

$$\mathcal{R}(n_j)(-\mathbf{a}) > \mathcal{R}(n_k)(-\mathbf{a}).$$

We will study the first case, the second one goes analogously. We claim it implies that

$$p_i + n_l \succ f_{\mathbf{a},b} \quad (7.7)$$

for any $n_l \in N_{I_-^{\sigma'', \sigma'}}$, which would be a contradiction since $f_{\mathbf{a},b} \succeq P \diamond N$.

Indeed, assume Equation (7.6) was true and pick any $n_l \in N_{I_-^{\sigma'', \sigma'}}$. Then, since $\zeta \subseteq f_{\mathbf{a},b}$, the following two equalities hold (the first one is just Equation (7.6)):

$$\begin{cases} \mathcal{R}(p_i)(-\mathbf{a}) > \mathcal{R}(p_k)(-\mathbf{a}) \\ p_k + n_l \in f_{\mathbf{a},b}. \end{cases}$$

By Proposition 3.2.4 and linearity of \mathcal{R} , this system can be re-written as

$$\begin{cases} \mathcal{R}(p_i)(-\mathbf{a}) > \mathcal{R}(p_k)(-\mathbf{a}) \\ \mathcal{R}(p_k) + \mathcal{R}(n_l) \ni (-\mathbf{a}, b), \end{cases}$$

that is,

$$\begin{cases} \mathcal{R}(p_i)(-\mathbf{a}) > \mathcal{R}(p_k)(-\mathbf{a}) \\ \mathcal{R}(p_k)(-\mathbf{a}) + \mathcal{R}(n_l)(-\mathbf{a}) = b. \end{cases}$$

We conclude that

$$\mathcal{R}(p_i)(-\mathbf{a}) + \mathcal{R}(n_l)(-\mathbf{a}) > b,$$

which again can be re-written as

$$\mathcal{R}(p_i) + \mathcal{R}(n_l) \succ (-\mathbf{a}, b) = \check{\mathcal{R}}(f_{\mathbf{a},b}).$$

Equation (7.7) then follows from Proposition 3.2.4.2. and linearity of \mathcal{R} . This finishes the proof of the first implication of the Proposition.

As in the proof of Proposition 6.1.6, reversing the above argument shows the other direction. \square

The following theorem makes precise the relationship between cells and faces that was introduced in the previous two Propositions:

Theorem 7.0.5 (Related to [18, Proposition 22]). *There exists a one-to-one correspondence between k -cells in $\mathcal{T}(P, N)$ and $(d - k)$ -cells in $\mathcal{U}(P \diamond N)$. Specifically, the following map is a bijection:*

$$\begin{aligned} \Psi: \mathcal{T}_k(P, N) &\xrightarrow{\sim} \mathcal{U}_{d-k}(P \diamond N) \\ \sigma = \sigma' \cap \sigma'' &\mapsto \mathcal{C}(P_{I_{\sigma', \sigma''}} \diamond N_{I_{\sigma'', \sigma'}}). \end{aligned}$$

Proof. Follows from Proposition 7.0.3 and Proposition 7.0.4 similarly to how Theorem 6.1.7 follows from Proposition 6.1.4 and Proposition 6.1.6. \square

7.1 Counting Affine Regions

In the special case where $\dim \sigma = d$, the result in Theorem 7.0.5 allows counting the number of affine regions defined by a ReLU network. Indeed, we will show in this section how the affine regions can be constructed from $\mathcal{T}_d(P, N)$ as a set of equivalence classes. We will then translate this observation to dual space.

Let $\sigma \in \mathcal{T}_d(P, N)$ be a d -cell in the tessellation induced by $\mathcal{N} := \mathcal{Q}(P) - \mathcal{Q}(N)$. Then \mathcal{N} is an affine map when restricted to σ . In particular, there exist a $p^\sigma \in P$ and an $n^\sigma \in N$ such that

$$\mathcal{N}(\mathbf{x}) = (\mathcal{R}(p^\sigma) - \mathcal{R}(n^\sigma))(\mathbf{x}) = \mathcal{R}(p^\sigma - n^\sigma) \quad \forall \mathbf{x} \in \sigma. \quad (7.8)$$

However, two d -cells can define the same affine map:

Proposition 7.1.1. *Let $\sigma, \sigma' \in \mathcal{T}_d(P, N)$ be two distinct d -cells. Then σ and σ' define the same affine map if and only if the corresponding vertices $\Psi(\sigma) = p^\sigma + n^\sigma$, $\Psi(\sigma') = p^{\sigma'} + n^{\sigma'} \in \mathcal{U}^*(P \diamond N)$ satisfy $p^\sigma - n^\sigma = p^{\sigma'} - n^{\sigma'}$.*

Proof. Follows directly from the definitions and Equation (7.8). \square

If two neighboring d -cells, i.e., two d -cells which share a $(d - 1)$ -face, define the same affine map, they are part of the same affine region. This implies that affine regions are more coarse than $\mathcal{T}_d(P, N)$. The rest of this section makes this observation more precise and translates it to dual space.

Definition 7.1.2 (Adjacency). We make the following two definitions:

- i) We say that two d -cells in $\mathcal{T}_d(P, N)$ are *adjacent* if they share a $(d - 1)$ -face.
- ii) We say two vertices $p_1 + n_1, p_2 + n_2 \in \mathcal{U}^*(P \diamond N)$ are *adjacent* if there exists an edge $\zeta \in \mathcal{U}_1(P \diamond N)$ going from $p_1 + n_1$ to $p_2 + n_2$.

The following proposition relates these two notions of adjacency:

Proposition 7.1.3. *Let $\sigma_1, \sigma_2 \in \mathcal{T}_d(P, N)$ be two distinct d -cells. Then σ and σ' are adjacent if and only if the corresponding vertices $\Psi(\sigma_1) = p_1 + n_1$, $\Psi(\sigma_2) = p_2 + n_2 \in \mathcal{U}^*(P \diamond N)$ are adjacent (that is, satisfy $p_1 - n_1 = p_2 - n_2$).*

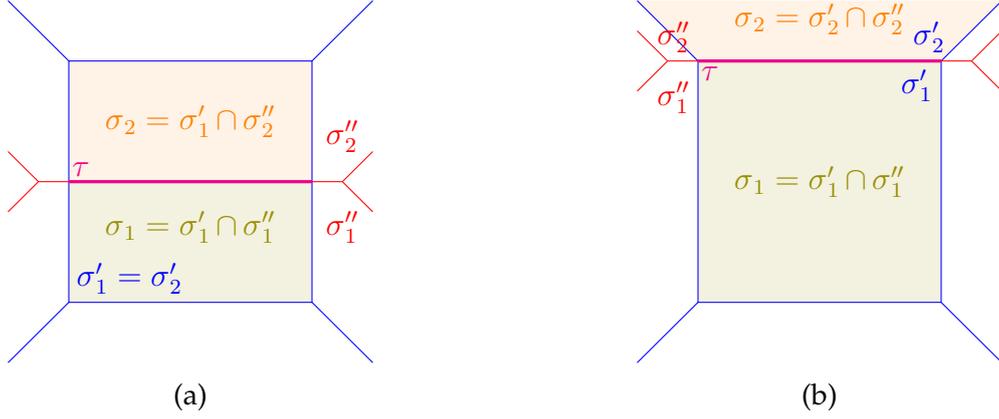


Figure 7.3: Illustration of the two cases in the proof of Proposition 7.1.3. The blue tessellation corresponds to $\mathcal{T}(P)$, the red tessellation corresponds to $\mathcal{T}(N)$. The shaded regions are the d -cells σ_1 and σ_2 . Subfigure (a) shows the first case, where $\sigma'_1 = \sigma'_2$ but $\sigma''_1 \neq \sigma''_2$. Subfigure (b) shows the second case, where $\sigma'_1 \neq \sigma'_2$ and $\sigma''_1 \neq \sigma''_2$. In both cases, the $(d-1)$ -cell τ is a face of both σ_1 and σ_2 .

Proof. “ \Rightarrow ”: Let $\sigma_1 = \sigma'_1 \cap \sigma''_1$ and $\sigma_2 = \sigma'_2 \cap \sigma''_2$ with $\sigma'_1, \sigma'_2 \in \mathcal{T}_d(P)$ and $\sigma''_1, \sigma''_2 \in \mathcal{T}_d(N)$ be adjacent.

Then σ_1 and σ_2 share a $(d-1)$ -cell $\tau \in \mathcal{T}_{d-1}(P, N)$.

We claim that, in the sense of Theorem 7.0.5, τ corresponds to a face

$$\Psi(\tau) = \mathcal{C}(\{p^{\sigma_1} + n^{\sigma_1}, p^{\sigma_2} + n^{\sigma_2}\}). \quad (7.9)$$

To see that this is true, differentiate two cases (see also Figure 7.3).

If $p^{\sigma_1} = p^{\sigma_2}$, i.e., σ_1 and σ_2 are defined by the same d -cell $\sigma'_1 = \sigma'_2 \in \mathcal{T}_d(P)$, then Equation (7.9) follows directly from the fact that there is only one implicit equality constraint $P_{\sigma'_1, \sigma'_2}^{\sigma'_1, \sigma'_2} = \{p^{\sigma_1}\}$ coming from P and Theorem 7.0.5.

We may thus assume that $p^{\sigma_1} \neq p^{\sigma_2}$ and $n^{\sigma_1} \neq n^{\sigma_2}$. In this case, there are four implicit equality constraint and

$$\Psi(\tau) = \mathcal{C}(\{p^{\sigma_1} + n^{\sigma_1}, p^{\sigma_1} + n^{\sigma_2}, p^{\sigma_2} + n^{\sigma_1}, p^{\sigma_2} + n^{\sigma_2}\}) \in \mathcal{U}(P \diamond N)$$

We claim that $p^{\sigma_1} + n^{\sigma_2}, p^{\sigma_2} + n^{\sigma_1} \notin \mathcal{U}^*(P \diamond N)$, which would show Equation (7.9). Indeed, let \mathcal{H} be the hyperplane defined by τ . Assume towards a contradiction that $p^{\sigma_1} + n^{\sigma_2} \in \mathcal{U}^*(P \diamond N)$. Then, by Theorem 7.0.5, $\sigma'_1 \cap \sigma''_2$ would be a d -cell in $\mathcal{T}_d(P, N)$. This is a contradiction since σ'_1 and σ''_2 lie on different sides of \mathcal{H} . The same argument shows that $p^{\sigma_2} + n^{\sigma_1} \notin \mathcal{U}^*(P \diamond N)$.

“ \Leftarrow ”: Assume $p_1 + n_1, p_2 + n_2 \in \mathcal{U}^*(P \diamond N)$ are adjacent. Let $\sigma_1 = \sigma'_1 \cap \sigma''_1, \sigma_2 = \sigma'_2 \cap \sigma''_2 \in \mathcal{T}_d(P, N)$ be the two d -cells related to $p_1 + n_1, p_2 + n_2$ and $\tau \in \mathcal{T}_{d-1}(P, N)$ the $(d-1)$ -cell related to the 1-face $\mathcal{C}(\{p_1 + n_1, p_2 + n_2\})$.

We claim that τ is a face of both σ_1 and σ_2 , which would conclude the proof.

By Theorem 7.0.5, it suffices to show that

$$\mathcal{C}(P_{\underline{=}}^{\sigma'_1, \sigma'_2} \diamond N_{\underline{=}}^{\sigma''_1, \sigma''_2}) = \mathcal{C}(\{p_1 + n_1, p_2 + n_2\}).$$

But this follows analogously to Equation (7.9). \square

We can now construct the equivalence relation which will identify adjacent d -cells in $\mathcal{T}_d(P, N)$ defining the same affine function:

Definition 7.1.4 (Path of d -cells). A *path of d -cells* is a sequence $(\sigma_1, \dots, \sigma_n) \subseteq \mathcal{T}_d(P, N)$ of d -cells such that

- i) $\mathcal{Q}(P) - \mathcal{Q}(N)$ defines the same affine map on σ_i and σ_{i+1} for all $i = 1, \dots, n - 1$
- ii) σ_i is adjacent to σ_{i+1} for all $i = 1, \dots, n - 1$.

We write $\mathfrak{P}(P, N)$ for the set of all paths of d -cells in $\mathcal{T}_d(P, N)$.

Definition 7.1.5 (Equivalence of d -cells). Given two d -cells $\sigma, \sigma' \in \mathcal{T}_d(P, N)$, we write $\sigma \sim \sigma'$ if there exists a path of d -cells from σ to σ' .

Clearly, \sim defines an equivalence relation.

By Proposition 7.1.3 and Proposition 7.1.1, this equivalence-relation translates to dual-space. This motivates the following definition:

Definition 7.1.6 (Path of dual points). A *path of dual points* is a sequence $(p^{\sigma_1} + n^{\sigma_1}, \dots, p^{\sigma_n} + n^{\sigma_n}) \subseteq \mathcal{U}^*(P \diamond N)$ of dual points such that

- i) $p^{\sigma_i} + n^{\sigma_i}$ is adjacent to $p^{\sigma_{i+1}} + n^{\sigma_{i+1}}$ for all $i = 1, \dots, n - 1$
- ii) $p^{\sigma_i} - n^{\sigma_i} = p^{\sigma_{i+1}} - n^{\sigma_{i+1}}$ for all $i = 1, \dots, n - 1$.

We write $\mathfrak{P}(P \diamond N)$ for the set of all paths of dual points in $\mathcal{U}^*(P \diamond N)$.

In particular, this definition induces an equivalence-relation \sim on $\mathcal{U}^*(P \diamond N)$, where $p_1 + n_1 \sim p_2 + n_2$ if and only if there exists a path of dual points from $p_1 + n_1$ to $p_2 + n_2$.

The following proposition relates paths in $\mathcal{T}_d(P, N)$ to paths in $\mathcal{U}^*(P \diamond N)$:

Proposition 7.1.7. *There exists a one-to-one correspondence between paths of d -cells in $\mathcal{T}_d(P, N)$ and paths of dual points in $\mathcal{U}^*(P \diamond N)$. It is given by*

$$\begin{aligned} \Theta: \mathfrak{P}(P, N) &\rightarrow \mathfrak{P}(P \diamond N) \\ (\sigma_1, \dots, \sigma_n) &\mapsto (p^{\sigma_1} + n^{\sigma_1}, \dots, p^{\sigma_n} + n^{\sigma_n}). \end{aligned}$$

Proof. Follows from Proposition 7.1.3 and Proposition 7.1.1. \square

We finally make precise the one-to-one correspondence between affine regions and equivalence classes of d -cells in $\mathcal{T}_d(P, N)$:

Corollary 7.1.8. *There exists a one-to-one-correspondence between affine regions of a ReLU network $\mathcal{Q}(P) - \mathcal{Q}(N): \mathbb{R}^d \rightarrow \mathbb{R}$ and equivalence classes in $\mathcal{U}^*(P \diamond N)/\sim$.*

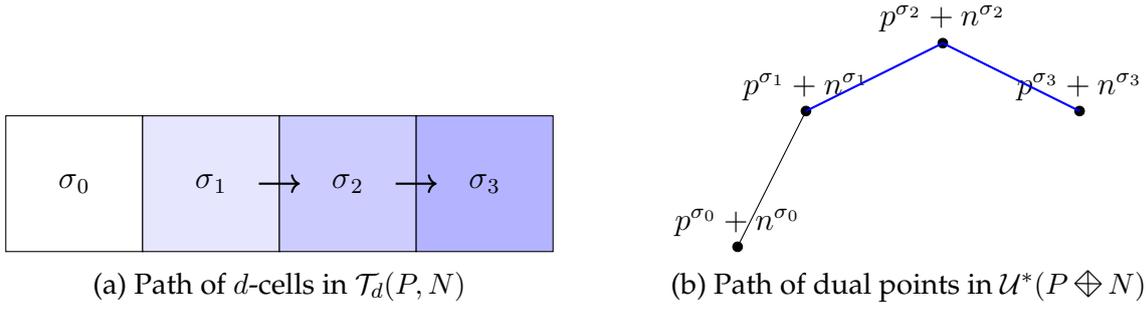


Figure 7.4: Example for the one-to-one correspondence between paths of d -cells in $\mathcal{T}_d(P, N)$ and paths of dual points in $\mathcal{U}^*(P \diamond N)$ (Proposition 7.1.7). Subfigure (a) shows a path of d -cells $(\sigma_1, \sigma_2, \sigma_3)$ (blue), consisting of adjacent cells that define the same affine function. In dual space (Subfigure b), this corresponds to a path (thick, blue) of adjacent vertices $p^{\sigma_i} + n^{\sigma_i}$ with the property that $p^{\sigma_i} - n^{\sigma_i} = p^{\sigma_j} - n^{\sigma_j}$ for all $i, j = 1, 2, 3$.

Proof. Follows from Theorem 7.0.5 and the fact that \sim identifies exactly adjacent cells that are part of the same affine region. \square

To better understand the space $\mathcal{U}^*(P \diamond N) / \sim$, define an unweighted graph $G = (V, E)$ with vertices $V := \mathcal{U}^*(P \diamond N) \subseteq \mathbb{R}^{d+1}$ and edges $E := \mathcal{U}_1(P \diamond N)$. If $d = 2$, then clearly G is planar.

The set $\mathcal{U}^*(P \diamond N) / \sim$ arises from the graph G by contracting exactly the paths in $\mathfrak{P}(P \diamond N)$ (see Figure 7.5).

Remark 7.1.9. We have seen above that d -cells in $\mathcal{T}(P, N)$ can be more fine than the affine regions of $\mathcal{N} = \mathcal{Q}(P) - \mathcal{Q}(N)$. Since also the activation regions [14, Definition 1] are generally finer than the affine regions [14, Lemma 3], one might ask whether the d -cells in $\mathcal{T}(P, N)$ are the same as the activation regions. However, this is not the case. As noted by Hanin and Rolnick [14, p.4], zeroing out a subnetwork may lead to different activation patterns that coalesce into a single linear region.

Importantly, the zeroed-out subnetwork does not affect the upper convex hulls of P and N and therefore does not influence the tessellation $\mathcal{T}(P, N)$. That is, two activation patterns that only differ in the zeroed-out subnetwork do not influence $\mathcal{T}(P, N)$.

Therefore, if \mathcal{N} restricts to the same affine map on two adjacent cells $\sigma_1, \sigma_2 \in \mathcal{T}_d(P, N)$, two adjacent activation regions which do not just differ by a zeroed-out subnetwork need to coalesce into the same affine region. We conjecture that this happens with probability zero.

By the above argument, this means that the corresponding points $p_1 + n_1$ and $p_2 + n_2$ in dual-space lie on the upper convex hull and satisfy $p_1 - n_1 = p_2 - n_2$. This furthermore motivates the conjecture, as this seems to be unlikely.

The considerations in Remark 7.1.9 lead us to conjecture that, in networks with random parameters, the d -cells are almost surely the same as the affine regions:

Conjecture 7.1.10. The number of affine regions of a random ReLU network $\mathcal{Q}(P) - \mathcal{Q}(N)$ is almost surely equal to the number of points in $\mathcal{U}^*(P \diamond N)$.

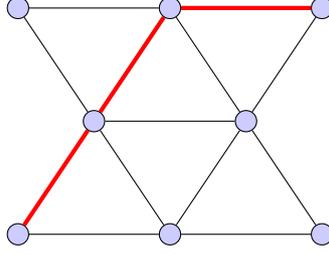


Figure 7.5: An example of the graph G induced by the 1-skeleton $\mathcal{U}_1(P \diamond N)$ for input-dimension $d = 2$. By Corollary 7.1.8, the affine regions induced by $\mathcal{Q}(P) - \mathcal{Q}(N)$ correspond to the vertices in the (multi)-graph G' obtained from G by identifying all the vertices along paths $P \in \mathfrak{P}(P \diamond N)$ (red).

Remark 7.1.11. The argument put forward in this chapter can be generalized to networks $\mathcal{N}: \mathbb{R}^d \rightarrow \mathbb{R}^o$ with output dimension $o > 1$. We leave this for future work. On a high level, the above construction applies to every coordinate function $f_i, i = 1, \dots, o$, of the network. In the end, one has to study intersections of cells in each coordinate to understand the tessellation induced by \mathcal{N} .

Example 7.1.12. In this subsection, we continue the example of the toy network $\mathcal{N} = \mathcal{Q}(P) - \mathcal{Q}(N)$ from Example 5.2.6 and Example 6.2.5. By Corollary 7.1.8, the number of affine regions defined by \mathcal{N} corresponds to the number of vertices in $\mathcal{U}(P \diamond N) / \sim$.

Specifically, we compute

$$\begin{aligned} P \diamond N = \{ & (7, 35, 10), (12, 34, 0), (4, 32, 8), (12, 34, 3), (3, 28, 11), (17, 24, 0), (13, 23, 2), \\ & (11, 36, 1), (11, 36, 4), (14, 21, 3), (2, 30, 12), (16, 26, 1), (16, 26, -1), (10, 38, 7), \\ & (10, 38, 4), (12, 19, 5), (6, 31, 8), (10, 38, 5), (6, 31, 11), (8, 33, 4), (11, 21, 6), \\ & (8, 33, 7), (4, 32, 10), (8, 36, 6), (5, 33, 9), (6, 28, 6), (8, 36, 9), (13, 23, 4), (11, 36, 3), \\ & (18, 22, -3), (5, 33, 12), (11, 36, 6), (5, 30, 7), (15, 22, 2), (7, 35, 5), (9, 34, 5), \\ & (7, 35, 8), (9, 34, 8), (8, 33, 6), (14, 24, 3), (15, 19, 0), (8, 33, 9), (9, 31, 3), \\ & (17, 24, -2), (9, 31, 6), (14, 21, 1), (10, 38, 2), (5, 30, 9), (7, 35, 7) \} \end{aligned}$$

and

$$\begin{aligned} \mathcal{U}^*(P \diamond N) = \{ & (18, 22, -3), (15, 19, 0), (10, 38, 7), (12, 19, 5), \\ & (17, 24, 0), (2, 30, 12), (3, 28, 11), (5, 33, 12) \} \end{aligned}$$

One can quickly see that $\mathfrak{P}(P \diamond N) = \emptyset$, i.e., there are no adjacent dual points $p_1 + n_1, p_2 + n_2 \in \mathcal{U}^*(P \diamond N)$ satisfying $p_1 - n_1 = p_2 - n_2$.

This implies $\mathcal{U}^*(P \diamond N) / \sim = \mathcal{U}^*(P \diamond N) = 8$, and \mathcal{N} thus defines 8 affine regions. This is confirmed by Figure 5.1a, which plots the tessellation induced by \mathcal{N} and contains 8 affine regions.

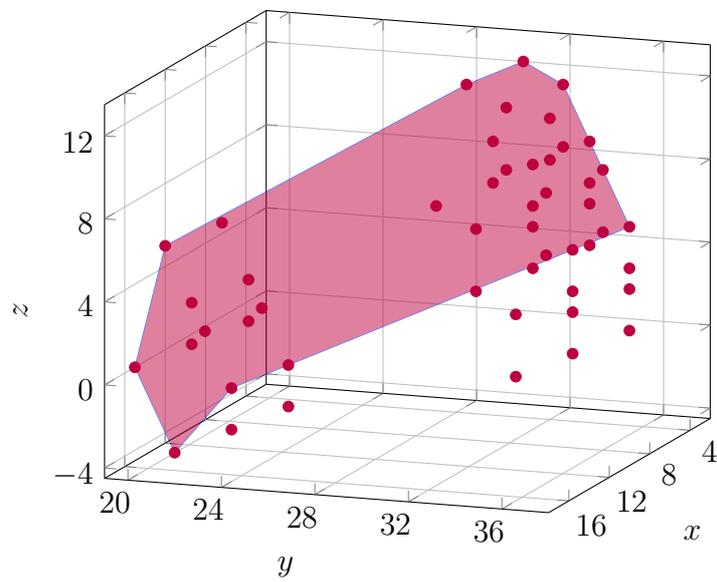


Figure 7.6: Two-dimensional toy-example defined in Equation (5.12). Purple points correspond to $P \diamond N$. The purple polygon is $\mathcal{U}(P \diamond N)$.

Chapter 8

Telgarsky's Sawtooth Network

We use this chapter to familiarize ourselves with the dual representation by specifically constructing it for a known network.

In his seminal work [15], Telgarsky constructed a binary classification problem for 2^k points which every shallow network with fewer than exponentially many nodes ($n \ll 2^k$) can not classify. In this chapter, we study this construction in the dual setting.

We furthermore present numerical experiments which provide evidence for the volume hypothesis.

8.1 Problem Setting

We start by introducing Telgarsky's main results [15]. Define $\mathfrak{R}(\rho_0; m, L)$ to be the class of neural networks with L layers, at most m nodes in each layer, output-dimension 1 and ReLU activations. Given a network $\mathcal{N} \in \mathfrak{R}(\rho_0; m, L)$, we denote by $\tilde{\mathcal{N}}$ the induced binary classifier and by $\mathcal{L}(\mathcal{N})$ the classification error of $\tilde{\mathcal{N}}$, given a dataset (for more details, see [15]).

The main theorem reads:

Theorem 8.1.1 ([15, Theorem 1.1]). *Let k be a positive integer, L the number of layers, m the number of nodes per layer with $m \leq 2^{(k-3)/(L-1)}$. Then there exists a collection of $n := 2^k$ points $((x_i, y_i))_{i=1, \dots, n}$ with $x_i \in [0, 1]$ and $y_i \in \{0, 1\}$ such that*

$$\min_{\mathcal{N} \in \mathfrak{R}(\rho_0; 2, 2k)} \mathcal{L}(\mathcal{N}) = 0 \quad \text{and} \quad \min_{\mathcal{M} \in \mathfrak{R}(\rho_0; m, L)} \mathcal{L}(\mathcal{M}) \geq \frac{1}{6}.$$

Telgarsky's proof is constructive. The collection of points is the k -ap:

Definition 8.1.2 (The k -ap). The k -ap (k -alternating-point problem) is the set of 2^k uniformly spaced points within $[0, 1 - 2^{-k}]$ with alternating labels. That is, $((x_i, y_i))_{i=1, \dots, k}$ with $x_i = i2^{-k}$ and $y_i = 0$ if i is even and $y_i = 1$ otherwise.

To construct a narrow network that can correctly classify the k -ap, he introduces the mirror-map:

Definition 8.1.3 (Mirror-Map). The *mirror-map* is defined as

$$f_m: \mathbb{R} \rightarrow \mathbb{R}$$

$$x \mapsto \begin{cases} 2x & 0 \leq x \leq \frac{1}{2} \\ 2(1-x) & \frac{1}{2} \leq x \leq 1 \\ 0 & \text{otherwise.} \end{cases}$$

This mirror-map can be represented as a ReLU network:

Lemma 8.1.4. *The mirror-map f_m can be written as*

$$f_m(x) = \mathfrak{f}_m(x)$$

where $\mathfrak{f}_m(x) \in \mathfrak{R}(\rho_0; 2, 2)$ is a neural network given by

$$\mathfrak{f}_m(x) = \rho_0(\mathbf{W}_2 \cdot \rho_0(\mathbf{W}_1 \cdot x + \mathbf{b}_1))$$

with

$$\mathbf{W}_1 = \begin{pmatrix} 1 \\ 1 \end{pmatrix}, \quad \mathbf{b}_1 = \begin{pmatrix} 0 \\ -\frac{1}{2} \end{pmatrix}, \quad \mathbf{W}_2 = (2 \quad -4).$$

In order to show Theorem 8.1.1, Telgarsky actually proves the following, stronger statement:

Theorem 8.1.5 ([15, Theorem 1.2]). *Let $k, m, L \in \mathbb{N}$ be positive integers. Given a t -sawtooth¹ $s: \mathbb{R} \rightarrow \mathbb{R}$ and $n := 2^k$ points specified by the k -ap, then*

$$\min_{\mathcal{N} \in \mathfrak{R}(\rho_0; 2, 2; k)} \mathcal{L}(\mathcal{N}) = 0 \quad \text{and} \quad \min_{\mathcal{M} \in \mathfrak{R}(s; m, L)} \mathcal{L}(\mathcal{M}) \geq \frac{n - 4(tm)^k}{3n}.$$

The proof uses the following lemma for the lower bound, which shows that the mirror-map f_m concatenated with itself k -times² correctly classifies the k -ap:

Lemma 8.1.6. *On the 2^k points specified by the k -ap,*

$$\mathcal{L}(f_m^k) = 0.$$

Telgarsky's construction incentivizes the use of deep networks, since it shows that shallow networks with fewer than exponentially many (in k) nodes have an error of at least $1/6$ on the k -ap, whereas there exists a deep recurrent network with 2 nodes in each of the $2k$ layers achieving zero error.

¹a t -sawtooth is a piecewise affine function with t pieces

²The k -fold concatenation of f_m with itself is what gives this construction its name, since f_m^k looks like the teeth of a saw.

8.2 The Dual Representation

In this section, we study the dual representation of powers of f_m . By Lemma 8.1.4, f_m^k can be written as a recurrent neural network of depth $2k$ and width 2 for any $k \in \mathbb{N}$.

Given two non-negative integers $a \leq b$, we introduce the short-hand notation $[a : b] := \{a, a + 1, \dots, b - 1, b\}$.

The following proposition establishes a recursive formulation for the dual representation of f_m^k :

Proposition 8.2.1. *Let (P_{2k}, N_{2k}) denote the dual representation of f_m^k . Then the dual representation of f_m^{k+1} is given by*

$$\begin{aligned} N_{2(k+1)} &= (4N_{2k} \diamond 2N_{2k} \diamond \{0\}) \cup (4N_{2k} \diamond 2N_{2k} \boxplus -2 \diamond \{0\}) \cup (4P_{2k} \setminus N_{2k} \diamond 2N_{2k} \boxplus -2) \\ P_{2(k+1)} &= (4N_{2k} \diamond 2N_{2k} \diamond \{0\}) \cup (4N_{2k} \diamond 2P_{2k} \setminus N_{2k}) \cup N_{2(k+1)}. \end{aligned}$$

Proof. Since $f_m: \mathbb{R} \rightarrow \mathbb{R}$ has output-dimension equal to 1, the vectors N_{2k} and P_{2k} just contain one set of dual points each. We identify them with these sets.

By Lemma 8.1.4, computing f_m^{k+1} comes down to post-composing a 2-layer-network with a $2k$ -layer-network. After post-composing the first layer of f_m with f_m^k , by Proposition 5.1.7 the dual representation (N_{2k+1}, P_{2k+1}) takes the form

$$(N_{2k+1})_i = (\mathbf{W}_1^-)_{i1} P_{2k} \diamond (\mathbf{W}_1^+)_{i1} N_{2k} \quad (8.1)$$

$$= \{0\} \diamond N_{2k} \quad (8.2)$$

and

$$(P_{2k+1})_i = (\mathbf{W}_1^+)_{i1} P_{2k} \diamond (\mathbf{W}_1^-)_{i1} N_{2k} \boxplus (\mathbf{b}_1)_i \cup (N_{2k+1})_i \quad (8.3)$$

$$= P_{2k} \boxplus (\mathbf{b}_1)_i \cup (\{0\} \diamond N_{2k}). \quad (8.4)$$

for $i = 1, 2$. Note that we keep $\diamond \{0\}$ in the above calculations since $N_0 = \emptyset$ and $\emptyset + \{0\} := \{0\} \neq \emptyset$ (corresponds to $k = 0$ in the proof of Corollary 5.1.8).

In order to compute the second post-composed layer, note that $N_{2k} \subseteq P_{2k}$ and thus

$$P_{2k} = N_{2k} \cup P_{2k} \setminus N_{2k}.$$

Furthermore, the set union distributes over the Minkowski sum s.t.

$$\begin{aligned} N_{2(k+1)} &= (\mathbf{W}_2^-)_{11} (P_{2k+1})_1 \diamond (\mathbf{W}_2^-)_{12} (P_{2k+1})_2 + (\mathbf{W}_2^+)_{11} (N_{2k+1})_1 \diamond (\mathbf{W}_2^+)_{12} (N_{2k+1})_2 \\ &= 4(P_{2k+1})_2 \diamond 2(N_{2k+1})_1 \\ &\stackrel{(8.3), (8.1)}{=} 4(P_{2k} \boxplus -\frac{1}{2} \cup \{0\} \diamond N_{2k}) \diamond 2(\{0\} \diamond N_{2k}) \\ &= (4P_{2k} \boxplus -2 \diamond 2N_{2k}) \cup (2N_{2k} \diamond 4N_{2k} \diamond \{0\}) \\ &= (4(N_{2k} \cup P_{2k} \setminus N_{2k}) \boxplus -2 \diamond 2N_{2k}) \cup (4N_{2k} \diamond 2N_{2k} \diamond \{0\}) \\ &= (4N_{2k} \diamond 2N_{2k} \diamond \{0\}) \cup (4N_{2k} \diamond 2N_{2k} \boxplus -2 \diamond \{0\}) \cup (4P_{2k} \setminus N_{2k} \diamond 2N_{2k} \boxplus -2) \end{aligned}$$

and

$$\begin{aligned}
P_{2(k+1)} &= (\mathbf{W}_2^+)_{11}(P_{2k+1})_1 \diamond (\mathbf{W}_2^+)_{12}(P_{2k+1})_2 \diamond (\mathbf{W}_2^-)_{11}(N_{2k+1})_1 \diamond (\mathbf{W}_2^-)_{12}(N_{2k+1})_2 \cup N_{2k} \\
&= 2(P_{2k+1})_1 \diamond 4(N_{2k+1})_2 \cup N_{2(k+1)} \\
&\stackrel{(8.3),(8.1)}{=} 2(P_{2k} \cup \{0 \diamond N_{2k}\} \diamond 4(\{0 \diamond N_{2k}\}) \cup N_{2(k+1)}) \\
&= (2N_{2k} \diamond 4N_{2k} \diamond \{0\}) \cup (4N_{2k} \diamond \{0\} \diamond 2P_{2k}) \cup N_{2(k+1)} \\
&= (2N_{2k} \diamond 4N_{2k} \diamond \{0\}) \cup (4N_{2k} \diamond 2N_{2k} \diamond \{0\}) \diamond (4N_{2k} \diamond 2P_{2k} \setminus N_{2k}) \cup N_{2(k+1)} \\
&= (4N_{2k} \diamond 2N_{2k} \diamond \{0\}) \cup (4N_{2k} \diamond 2P_{2k} \setminus N_{2k}) \cup N_{2(k+1)}.
\end{aligned}$$

This finishes the proof. \square

Remark 8.2.2. We note two observations related to Proposition 8.2.1:

- i) For $k \geq 1$, it holds that $P_{2k}, N_{2k} \neq \emptyset$ and one can drop the $\diamond \{0\}$.
- ii) The reader may wonder why we introduced $P_{2k} \setminus N_{2k}$ in the proof of the Proposition. However, while this artificial step complicates the recursive identities for $P_{2(k+1)}$ and $N_{2(k+1)}$, it will pay off later. Indeed, we will see that computing N_{2k} and $P_{2k} \setminus N_{2k}$ is easier than computing N_{2k} and P_{2k} directly.

Next, we define the following two quantities:

Definition 8.2.3. For any $j \in \mathbb{N}$, we define

$$\begin{aligned}
h_0^j &:= 2 \frac{6^j - 1}{5} \\
g_0^j &:= \frac{3^j - 1}{2}.
\end{aligned}$$

Remark 8.2.4. Interestingly, h_0^j and g_0^j are the values that powers of certain functions take at zero. Indeed, define the two helper-functions $f: \mathbb{R} \rightarrow \mathbb{N}$, $h(x) = 6x + 2$ and $g: \mathbb{R} \rightarrow \mathbb{N}$, $g(x) = 3x + 1$. Then

$$\begin{aligned}
h_0^j &= 2 \frac{6^j - 1}{5} = \sum_{i=0}^{j-1} 2 \cdot 6^i \\
g_0^j &= \frac{3^j - 1}{2} = \sum_{i=0}^{j-1} 3^i.
\end{aligned}$$

The following are useful helping Lemmas:

Lemma 8.2.5. *The following equations of sets hold:*

$$[0 : 3^k] = [0 : 3^{k-1}] \diamond 2 [0 : 3^{k-1}] \quad (8.5)$$

$$[1 : 3^k - 1] = [0 : 3^{k-1}] \diamond 4 [0 : g_0^{k-1}] + 1 \quad (8.6)$$

$$[0 : g_0^k] = [0 : 3^{k-1}] \diamond [0 : g_0^{k-1}] \quad (8.7)$$

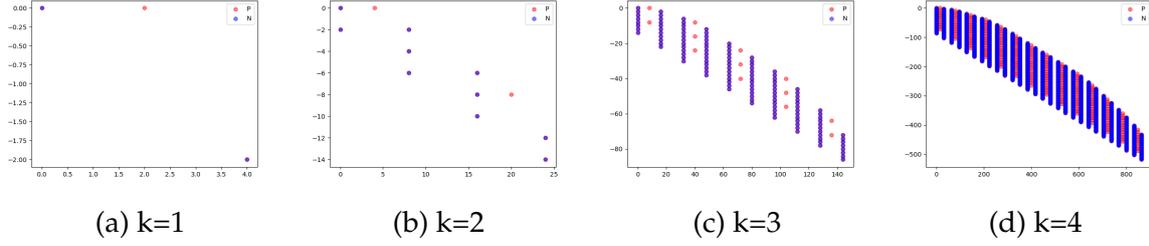


Figure 8.1: Unreduced representation of the sawtooth function f_m^k for $k = 1, 2, 3, 4$ (points in P are red and points in N are blue). One vertical “pillar” corresponds to the inner union in Proposition 8.2.6. The horizontal structure, i.e. the lining up of the vertical pillars, corresponds to the outer union.

Proof. The proof is straight forward. We still provide it for completeness, starting with Equation (8.5). Note that

$$2 [0 : 3^{k-1}] \diamond \{0, 1\} = [0 : 2 \cdot 3^{k-1} + 1]$$

and thus

$$2 [0 : 3^{k-1}] \diamond [0 : 3^{k-1}] = [0 : 2 \cdot 3^{k-1} + 3^{k-1}] = [0 : 3^k].$$

For Equation (8.6), proceed analogously:

$$4[0 : g_0^{k-1}] \diamond \{0, 1, 2, 3\} = [0 : 4g_0^{k-1} + 3] = [0 : 2(3^{k-1} - 1) + 3]$$

and thus

$$[0 : 3^{k-1}] \diamond 4[0 : g_0^{k-1}] = [0 : 3^{k-1} + 2(3^{k-1} - 1)] = [0 : 3^k - 2].$$

Finally, for Equation (8.7),

$$[0 : 3^{k-1}] \diamond [0 : g_0^{k-1}] = \left[0 : 3^{k-1} + \frac{3^{k-1} - 1}{2}\right] = [0 : g_0^k].$$

This concludes the proof. \square

With these helping Lemmas at hand, the following proposition provides a recursive formulation for the dual representation of the sawtooth function:

Proposition 8.2.6. *Assume that the dual representation of f_m^k with $l \geq 2$ can be written as*

$$N_{2k} = \bigcup_{j=0}^{3^{k-1}} \bigcup_{a \in S(k,j)} \{(j2^{k+1}, -h_0^{k-1} - j2^k + 2a)\} \quad (8.8)$$

and

$$P_{2k} = N_{2k} \sqcup H_{2k},$$

where

$$H_{2k} = \bigcup_{j=0}^{g_0^{k-1}} \bigcup_{b \in Z(k,j)} \{(2^k + j2^{k+2}, 2^k - 2 - h_0^{k-1} - j2^{k+1} + 8b)\} \quad (8.9)$$

for some sets $S(k, j), Z(k, j) \subseteq \mathbb{N}$. Then Equation (8.8) and Equation (8.9) also hold for $k + 1$. The sets $S(k + 1, q), q \in [0 : 3^k]$, are given by

$$S(k + 1, q) = \begin{cases} S'(k + 1, q), & q \in \{0, 3^k\} \\ S'(k + 1, q) + S''(k + 1, q) & \text{o.w.}, \end{cases}$$

where

$$S'(k + 1, q) := \{4o + 2p + \delta \mid \delta \in \{0, 1\}, \exists i, j \in [0 : 3^{k-1}] \text{ s.t. } q = 2i + j, \\ o \in S(k, i), p \in S(k, j)\}$$

and

$$S''(k + 1, q) := \{3 \cdot 2^k - 4 + 2o + 16p \mid \exists i \in [0 : 3^{k-1}], j \in [0 : 3^{k-1}] \text{ s.t. } q = i + 4j + 1, \\ o \in S(k, i), p \in Z(k, j)\}$$

The sets $Z(k + 1, q), q \in [0 : 3^k]$, are given by

$$Z(k + 1, q) = \{o + 2p \mid \exists i \in [0 : 3^{k-1}], j \in [0 : 3^{k-1}] \text{ s.t. } q = i + j, o \in S(k, i), p \in Z(k, j)\}.$$

Proof. Since $N_{2k} \subseteq P_{2k}$, note first that $H_{2k} = P_{2k} \setminus N_{2k}$.

We start by showing that Equation (8.8) also holds for $k + 1$ with the defined $S(k + 1, \cdot)$.

By Proposition 8.2.1,

$$N_{2(k+1)} = A \cup B \cup C$$

where

$$\begin{aligned} A &:= 4N_{2k} \diamond 2N_{2k} \\ B &:= 4N_{2k} \diamond 2N_{2k} \boxplus -2 \\ C &:= 4P_{2k} \setminus N_{2k} \diamond 2N_{2k} \boxplus -2. \end{aligned}$$

We start by computing A :

$$\begin{aligned} A &= 4N_{2k} \diamond 2N_{2k} \\ &= \{(4i2^{k+1}, -4h_0^{k-1} - 4i2^k + 4 \cdot 2o) \mid i \in [0 : 3^{k-1}], o \in S(k, i)\} \\ &\quad \diamond \{(2j2^{k+1}, -2h_0^{k-1} - 2j2^k + 2 \cdot 2p) \mid j \in [0 : 3^{k-1}], p \in S(k, j)\} \\ &= \{(2^{k+2}(2i + j), -h_0^k - 2^{k+1}(2i + j) + 8o + 4p + 2) \mid \\ &\quad \mid i, j \in [0 : 3^{k-1}], o \in S(k, i), p \in S(k, j)\} \\ &= \{(2^{k+2}q, -h_0^k - 2^{k+1}q + 2(4o + 2p + 1)) \mid \\ &\quad \mid \exists i, j \in [0 : 3^{k-1}] \text{ s.t. } q = 2i + j, o \in S(k, i), p \in S(k, j)\}. \end{aligned}$$

Since $B = A \boxplus -b$, it follows readily that

$$\begin{aligned} B &= \{(2^{k+2}q, -h_0^k - 2^{k+1}q + 2(4o + 2p)) \mid \\ &\quad \mid \exists i, j \in [0 : 3^{k-1}] \text{ s.t. } q = 2i + j, o \in S(k, i), p \in S(k, j)\}, \end{aligned}$$

from which we conclude, using Equation (8.5) and the definition of $S'(k+1, q)$, that

$$A \cup B = \{(2^{k+2}q, -h_0^k - 2^{k+1}q + 2r \mid q \in [0 : 3^k], r \in S'(k+1, q)) \mid \}.$$

Finally,

$$\begin{aligned} C &= 4P_{2k} \setminus N_{2k} \diamond 2N_{2k} \boxplus -2 \\ &= \{(4 \cdot 2^k + 4j2^{k+2}, 4 \cdot 2^k - 8 - 4h_0^{k-1} - 4j2^{k+1} + 32p - 2) \mid j \in [0 : g_0^{k-1}], p \in Z(k, j)\} \\ &\quad + \{(2i2^{k+1}, -2h_0^{k-1} - 2i2^k + 4o) \mid i \in [0 : 3^{k-1}], o \in S(k, i)\} \\ &= \{(2^{k+2}(i+4j+1), -h_0^k - 2^{k+1}(i+4j+1) + 2^{k+2} + 2^{k+1} - 8 + 32p + 4o) \mid \\ &\quad \mid i \in [0 : 3^{k-1}], o \in S(k, i), j \in [0 : g_0^{k-1}], p \in Z(k, j)\} \\ &= \{(2^{k+2}(i+4j+1), -h_0^k - 2^{k+1}(i+4j+1) + 3 \cdot 2^{k+1} - 8 + 32p + 4o) \mid \\ &\quad \mid i \in [0 : 3^{k-1}], j \in [0 : g_0^{k-1}], o \in S(k, i), p \in Z(k, j)\} \\ &= \{(2^{k+2}q, -h_0^k - 2^{k+1}q + 2(3 \cdot 2^k - 4 + 16p + 2k)) \mid \\ &\quad \mid \exists i \in [0 : 3^{k-1}], j \in [0 : g_0^{k-1}] \text{ s.t. } q = i + 4j + 1, o \in S(k, i), p \in Z(k, j)\} \\ &= \{(2^{k+2}q, -h_0^k - 2^{k+1}q + 2r) \mid q \in [1 : 3^k - 1], r \in S''(k+1, q)\}, \end{aligned}$$

where in last step we used Equation (8.6) and the definition of $S''(k+1, q)$.

Putting together these results for $A \cup B$ and C shows that Equation (8.8) also holds for N_{2k} with $S(k+1, \cdot)$ defined as in the Proposition.

It remains to show that Equation (8.9) holds for $k+1$ as well. Note again that, since $N_{2(k+1)} \subseteq P_{2(k+1)}$, we can again identify $H_{2(k+1)} = P_{2(k+1)} \setminus N_{2(k+1)}$. By Proposition 8.2.1, this implies that

$$H_{2(k+1)} = H'_{2(k+1)} \setminus N_{2(k+1)}, \quad \text{where} \quad H'_{2(k+1)} := 4N_{2k} \diamond 2P_{2k} \setminus N_{2k}.$$

As a first step, we compute $H'_{2(k+1)}$ to be

$$\begin{aligned} H'_{2(k+1)} &= 4N_{2k} \diamond 2P_{2k} \setminus N_{2k} \\ &= \{(4i2^{k+1}, -4h_0^{k-1} - 4i2^k + 8o) \mid i \in [0 : 3^{k-1}], o \in S(k, i)\} \\ &\quad \diamond \{(2 \cdot 2^k + 2j2^{k+1}, 2 \cdot 2^k - 4 - 2h_0^{k-1} - 2j2^{k+1} + 16p) \mid j \in [0 : g_0^{k-1}], p \in Z(k, j)\} \\ &= \{(2^{k+1} + 2^{k+3}(i+j), 2^{k+1} - 2 - h_0^k - 2^{k+2}(i+j) + 8o + 16p) \mid \\ &\quad \mid i \in [0 : 3^{k-1}], j \in [0 : g_0^{k-1}], o \in S(k, i), p \in Z(k, j)\} \\ &= \{(2^{k+1} + 2^{k+3}q, 2^{k+1} - 2 - h_0^k - 2^{k+2}q + 8(o+2p)) \mid \\ &\quad \mid \exists i \in [0 : 3^{k-1}], j \in [0 : g_0^{k-1}] \text{ s.t. } q = i + j, o \in S(k, i), p \in Z(k, j)\} \\ &= \{(2^{k+1} + 2^{k+3}q, 2^{k+1} - 2 - h_0^k - 2^{k+2}q + 8r) \mid q \in [0 : g_0^k], r \in Z(k+1, q)\}, \end{aligned}$$

where in the last step we used Equation (8.7) and the definition of $Z(k+1, q)$.

Comparing this result with the above result for $N_{2(k+1)}$, we can see that $H'_{2(k+1)} \cap N_{2(k+1)} = \emptyset$ (the x -coordinates can't match) and thus

$$H'_{2(k+1)} = H_{2(k+1)}.$$

This shows Equation (8.9) also holds for $N+1$ and concludes the proof. \square

An example of the dual representation derived in Proposition 8.2.6 can be found in Figure 8.1.

We conjecture that $S(k, j)$ and $Z(k, j)$ are actually contiguous in \mathbb{N} :³

Conjecture 8.2.7. For every $k \geq 1$ and $j \in [0 : 3^{k-1}]$, let $s(k, j) := \max S(k, j)$ and $z(k, j) := \max Z(k, j)$. Then

$$\begin{aligned} S(k, j) &= [0 : s(k, j)] \\ Z(k, j) &= [0 : z(k, j)]. \end{aligned}$$

The following proposition starts the induction for Proposition 8.2.6:

Proposition 8.2.8. For $k = 2$, the sawtooth-function f_m^2 has dual representation as given in Proposition 8.2.6.

Proof. A simple calculation using Proposition 8.2.1 shows that

$$\begin{aligned} N_4 = \{ & (0, 0), (0, -2) \\ & (8, -2), (8, -4), (8, -6) \\ & (16, -6), (16, -8), (16, -10) \\ & (24, -12), (24, -14)\}. \end{aligned}$$

A simple calculation reveals that N_4 can consequently be written as

$$\begin{aligned} N_4 &= \bigcup_{j=0}^3 \bigcup_{a \in S(2, j)} \{(8j, -2 - 4j + 2a)\} \\ &= \bigcup_{j=0}^{3^{k-1}} \bigcup_{a \in S(k, j)} \{(j2^{k+1}, -h_0^{k-1} - j2^k + 2a)\} \end{aligned}$$

where

$$\begin{aligned} S(2, 0) &= \{0, 1\} \\ S(2, 1) &= \{0, 1, 2\} \\ S(2, 2) &= \{0, 1, 2\} \\ S(2, 3) &= \{0, 1\}. \end{aligned}$$

The set N_4 is plotted in blue in Figure 8.1b.

Similarly, one can show using Proposition 8.2.1 that

$$P_4 \setminus N_4 = \{(4, 0), (20, -8)\},$$

³at least this seems to hold computationally. Remains to be shown.

which can be written as

$$\begin{aligned} P_4 \setminus N_4 &= \bigcup_{j=0}^1 \bigcup_{b \in Z(k,j)} \{(4 + 16j, -8j + 8b)\} \\ &= \bigcup_{j=0}^{g_0^{k-1}} \bigcup_{b \in Z(k,j)} \{(2^k + j2^{k+2}, 2^k - 2 - h_0^{k-1} - j2^{k+1} + 8b)\} \end{aligned}$$

where

$$\begin{aligned} Z(2, 0) &= \{0\} \\ Z(2, 1) &= \{0\}. \end{aligned}$$

The points in P_4 are plotted in red in Figure 8.1b. □

Restricting the sets P_{2k} and N_{2k} to their upper convex hulls, the following theorem provides the reduced dual representation of f_m^k :

Theorem 8.2.9. *For any integer $k \geq 2$, the sawtooth-function f_m^k can be represented as $f_m^k = \mathcal{Q}(P_{2k}) - \mathcal{Q}(N_{2k})$ with*

$$N_{2k} = \{(j2^{k+1}, -h_0^{k-1} - j2^k + 2s(k, j)) \mid j \in [0 : 3^{k-1}]\} \quad (8.10)$$

and

$$P_{2k} = N_{2k} \sqcup H_{2k},$$

where

$$H_{2k} = \{(2^k + j2^{k+2}, 2^k - 2 - h_0^{k-1} - j2^{k+1} + 8z(k, j)) \mid j \in [0 : g_0^{k-1}]\} \quad (8.11)$$

for some integers $s(k, j), z(k, j) \in \mathbb{N}$. We call this the reduced representation of f_m^k .

Proof. Follows from Proposition 8.2.6, Proposition 8.2.8 and Corollary 3.3.3. □

We conclude this section by pointing out how Theorem 8.2.9 supports Conjecture 5.1.20; both P_{2k} and N_{2k} are way bigger than $\mathcal{U}^*(P_{2k})$ and $\mathcal{U}^*(N_{2k})$, respectively (only the uppermost points of the vertical pillars from Proposition 8.2.6 lie in the upper convex hulls, see Figure 8.1).

8.3 Experiments

The careful design of the sawtooth function ensures that it has a high number of break-points. In this section, we study how the complexity of the decision boundary and the number of affine regions change if the carefully designed weights and biases are randomized.

Define f_m^{rand} to be the following random function:

$$\begin{aligned} f_m^{\text{rand}} : \mathbb{R} &\rightarrow \mathbb{R} \\ \mathbf{x} &\mapsto \mathbf{W}_2 \cdot \rho_0(\mathbf{W}_1 \cdot \mathbf{x} + \mathbf{b}_1) + \mathbf{b}_2 \end{aligned}$$

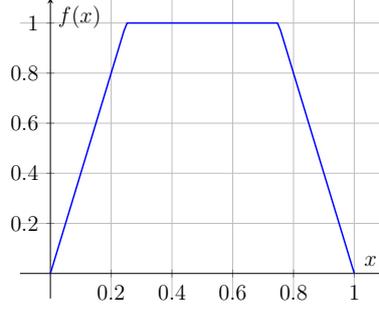


Figure 8.2: A "flattened cone", resulting from $W_1 = (-1, 1)^T$, $W_2 = (-4, -4)$, $b_1 = (1/4, -3/4)^T$, $b_2 = 1$.

for some random matrices $W_1 \in \mathbb{R}^{1,2}$, $W_2 \in \mathbb{R}^{2,1}$ with *i.i.d.* Gaussian entries and vectors $b_1 \in \mathbb{R}^2$, $b_2 \in \mathbb{R}$ with *i.i.d.* Gaussian entries. We write $f_m^{\text{rand}} \sim \mathfrak{N}^{\text{rand}}$ for a random function ("block") of that form.

Remark 8.3.1. The reader might notice that, compared to the deterministic case in Lemma 8.1.4, this definition is missing the outer nonlinearity. Indeed, for $n > 1$, we could also have discarded it in the deterministic case. However, we chose to stay close to the original construction in [15]. In particular, the reduced representation (Theorem 8.2.9) is the same with or without the outer nonlinearity. In the random case, we go without the nonlinearity in order to restrict the network's effective number of ReLU layers.

Definition 8.3.2. Given $a, b \in \mathbb{N}_0$, we define $\mathfrak{N}(a, b)$ to be the distribution of neural networks obtained by concatenating a consecutive deterministic blocks with b consecutive random blocks. That is,

$$\underbrace{f_{m,1}^{\text{rand}} \circ \dots \circ f_{m,b}^{\text{rand}}}_{b \text{ times}} \circ \underbrace{f_m \circ \dots \circ f_m}_{a \text{ times}} \sim \mathfrak{N}(a, b)$$

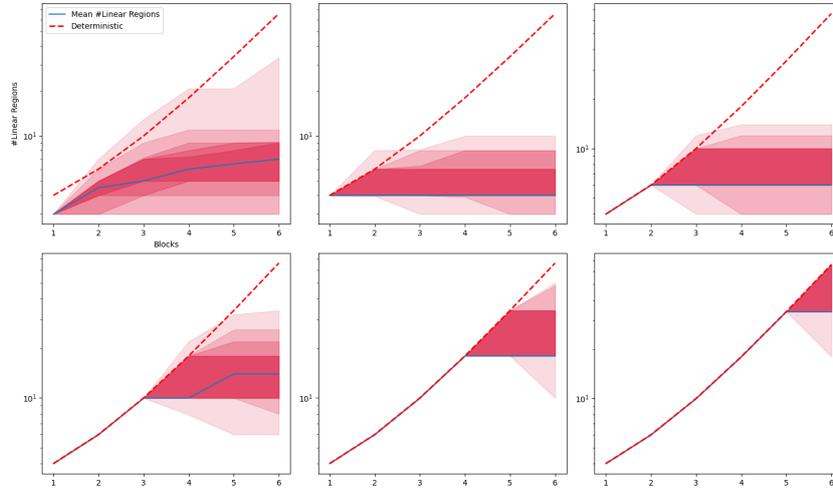
for $f_{m,1}^{\text{rand}}, \dots, f_{m,b}^{\text{rand}} \sim \mathfrak{N}^{\text{rand}}$.

For the following experiments, we use Conjecture 7.1.10 to count the number of affine regions defined by a network randomly drawn from $\mathfrak{N}(a, b)$, and Proposition 6.2.2 to assess its decision boundaries complexity.

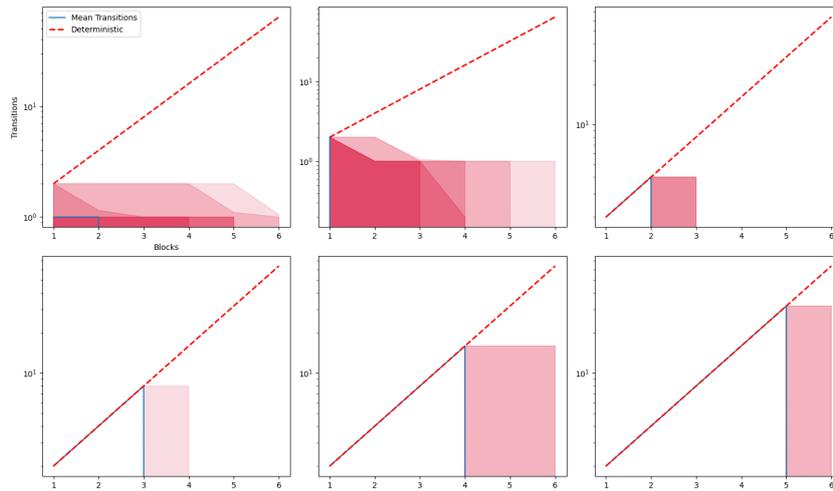
Figure 8.3a shows the behavior of both complexity measures, given a fixed network depth but a varying ratio of random blocks to deterministic blocks.

There, it is evident that both the complexity of the decision boundary and the number of affine regions grow more slowly or even decay compared to a purely deterministic sawtooth network. This provides evidence that the exponential complexity regime (introduced in Section 1.2) is unstable/sharp. In order to enter the regime, Telgarsky carefully constructed the weights and biases, enabling exponential complexity (in the number of layers).

As soon as the carefully constructed parameters are replaced by random ones, the network quickly transitions to the subexponential complexity regime with subexponential complexity (compare the blue line and the red dashed line in Figure 8.3a).



(a) Counting affine regions



(b) Counting Decision Boundary Pieces

Figure 8.3: Counting affine regions (Subfigure (a)) and transitions (Subfigure (b)) in different settings. Networks are made up of $k = 7$ blocks. Graphs in Subfigure (a) show $|\mathcal{U}^*(P_{2k} \diamond N_{2k})|$, i.e. the number of affine regions, after every block $l \leq k$, and graphs in Subfigure (b) count the number of edges between P_{2k} and N_{2k} in $\mathcal{U}(P_{2k} \cup N_{2k})$. Inside Subfigure, from the top left to the bottom right, the networks are first drawn from $\mathfrak{N}(0, 7)$, then from $\mathfrak{N}(1, 6)$ etc. The red dashed line indicates the number of affine regions (transitions) of a purely deterministic sawtooth network. Shaded regions indicate percentiles, starting from 70% and going to 95% in steps of 5%. The blue line is the empirical mean.

Remark 8.3.3. Note that, as shown in Figure 8.3a, it is possible, though unlikely, for a (partially) random network to have more affine regions than a purely deterministic sawtooth network. This can occur if all the breakpoints associated with the first random subnetwork fall within the interval $[0, 1]$. For instance, a random layer may resemble a "flattened cone" (see Figure 8.2), resulting in more affine regions than a regular cone in Telgarsky's construction.

Finally, Algorithm 2 and Algorithm 3 show how we compute the number of d -cells and the number of linear pieces in the decision boundary of a given ReLU network.

Algorithm 2 Counting Boundary Complexity

- 1: **Input:** Neural network $\mathcal{N} : \mathbb{R}^d \rightarrow \mathbb{R}$ with L layers, sample $\mathbf{x} \in \mathbb{R}^d$
 - 2: **Output:** Boundary complexity of \mathcal{N}
 - 3:
 - 4: **Initialize** $P_0 = (\{\mathbf{e}_1, 0\}, \dots, \{\mathbf{e}_d, 0\})$ and $N_0 = (\emptyset)$
 - 5: **for** layer $l = 1$ to L **do**
 - 6: Decompose \mathbf{W}_l into positive part \mathbf{W}_l^+ and negative part \mathbf{W}_l^-
 - 7: Compute $N_l = (\mathbf{W}_l^- P_{l-1}) \diamond (\mathbf{W}_l^+ N_{l-1})$
 - 8: Compute $P_l = ((\mathbf{W}_l^+ P_{l-1}) \diamond (\mathbf{W}_l^- N_{l-1})) \boxplus \mathbf{b}_l \cup (N_l \boxplus t_l)$
 - 9: **end for**
 - 10:
 - 11: Find the upper convex hull $\mathcal{U}(P_L \cup N_L)$
 - 12: **for** each 1-cell (edge) in $\mathcal{U}(P_L \cup N_L)$ **do**
 - 13: **if** the 1-cell joins a point in P_L to a point in N_L **then**
 - 14: Mark as a boundary piece
 - 15: **end if**
 - 16: **end for**
 - 17: **Return:** Total boundary piece count
-

Algorithm 3 Counting affine regions

- 1: **Input:** Neural network $\mathcal{N} : \mathbb{R}^d \rightarrow \mathbb{R}$ with L layers, sample $\mathbf{x} \in \mathbb{R}^d$
 - 2: **Output:** Number of d -cells defined by \mathcal{N}
 - 3:
 - 4: **Initialize** $P_0 = (\{\mathbf{e}_1, 0\}, \dots, \{\mathbf{e}_d, 0\})$ and $N_0 = (\emptyset)$
 - 5: **for** layer $l = 1$ to L **do**
 - 6: Decompose \mathbf{W}_l into positive part \mathbf{W}_l^+ and negative part \mathbf{W}_l^-
 - 7: Compute $N_l = (\mathbf{W}_l^- P_{l-1}) \diamond (\mathbf{W}_l^+ N_{l-1})$
 - 8: Compute $P_l = ((\mathbf{W}_l^+ P_{l-1}) \diamond (\mathbf{W}_l^- N_{l-1})) \boxplus \mathbf{b}_l \cup (N_l \boxplus t_l)$
 - 9: **end for**
 - 10:
 - 11: Find the upper convex hull $\mathcal{U}(P_L \diamond N_L)$
 - 12: **Return:** The number of vertices in $\mathcal{U}^*(P_L \diamond N_L)$.
-

Chapter 9

One Layer – Random Weights

Having familiarized ourselves with the dual representation by explicitly constructing it for the sawtooth function, we devote the remainder of this work to studying the volume hypothesis. We begin with a quick recap.

Chapter 5 established a dual representation $\mathcal{N} = \mathcal{Q}(P) - \mathcal{Q}(N)$ for every fully connected feedforward network \mathcal{N} with ReLU activations. The complexity of \mathcal{N} is closely tied to properties derived from upper convex hulls associated with its dual representation. By Corollary 7.1.8, the number of affine regions induced by \mathcal{N} equals the number of vertices in $\mathcal{U}(P \diamond N) / \sim$. Moreover, Proposition 6.2.2 states that the number of linear pieces within the decision boundary equals the number of edges in $\mathcal{U}(P \cup N)$ containing points from both P and from N .

In Chapter 8, we computed the dual representation of Telgarsky’s sawtooth network. Experiments confirmed the instability of the exponential complexity regime, as randomizing the last layers resulted in a sub-exponential number of affine regions. In this chapter, we analyze this phenomenon by explicitly computing the marginal gain in complexity achieved by appending a random ReLU layer to a deterministic network.

Since the dual representation of the sawtooth function is challenging to handle (see Theorem 8.2.9), we consider a slight simplification thereof.

Throughout this section, let $n \geq 2$ be an integer. We now introduce the deterministic subnetwork mathematically.

Let N_0 and P_0 be sets of $n + 1$ points each, defined by

$$N_0 = \left\{ \left(\cos \left(\frac{\pi}{2} \frac{2j}{2n+1} \right), \sin \left(\frac{\pi}{2} \frac{2j}{2n+1} \right) \right) \mid 0 \leq j \leq n \right\}$$
$$P_0 = \left\{ \left(\cos \left(\frac{\pi}{2} \frac{2j+1}{2n+1} \right), \sin \left(\frac{\pi}{2} \frac{2j+1}{2n+1} \right) \right) \mid 0 \leq j \leq n \right\}.$$

In particular, as in Telgarsky’s network, which inspired this construction, we conjecture that these points define a function $\mathbb{R} \rightarrow \mathbb{R}$ that is realized by a network $\mathcal{Q}(P_0) - \mathcal{Q}(N_0)$ in the exponential complexity regime.

Figure 9.1 shows the sets P_0 and N_0 for the case $n = 4$, as well as the DCPA function they define. As a prove of concept, note that the number of edges in $\mathcal{U}(P_0 \cup N_0)$ containing points from both N_0 and P_0 equals 9, which is also the the number of times the graph of the network (P_0, N_0) goes through zero (i.e., the number of pieces in its decision boundary). Similarly, there are 9 affine pieces, which, as we will see in Proposition 9.2.1, equals $|\mathcal{U}^*(P_0 \diamond N_0)|$.

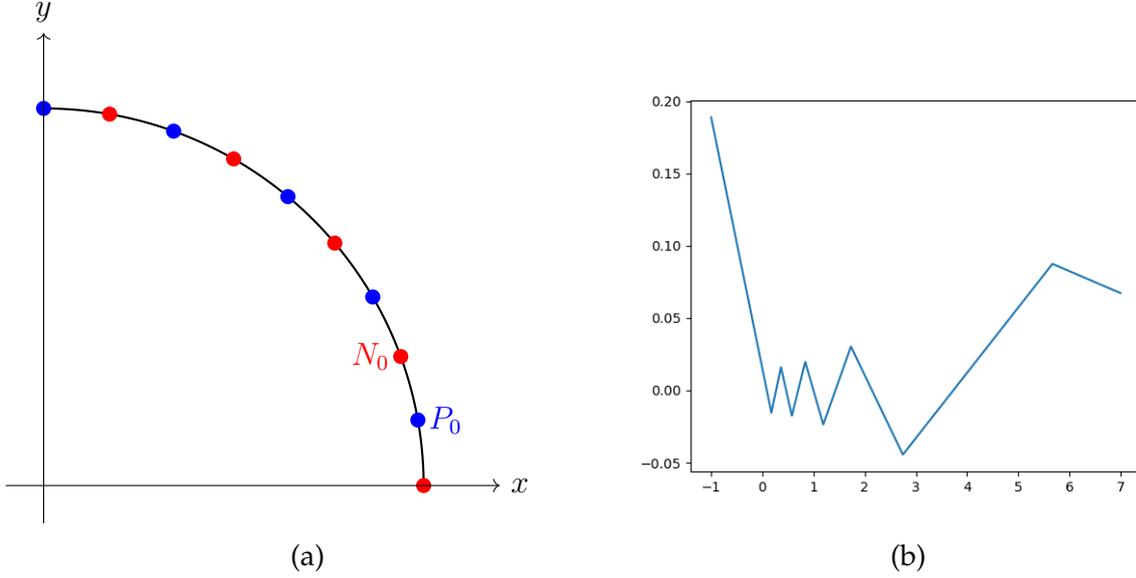


Figure 9.1: Circle construction for $n = 4$. Subfigure (a) shows the dual representation. Subfigure (b) plots the corresponding DCPA function.

The additional one-dimensional ReLU layer $x \mapsto \sigma(wx + b)$, where $w, b \sim \mathcal{N}(0, 1)$, transforms these points according to Proposition 5.1.7, resulting in an enlarged network with dual representation

$$N_1 = w^- P_0 \diamond w^+ N_0 \quad (9.1)$$

$$P_1 = (w^+ P_0 \diamond w^- N_0 \boxplus b) \cup N_1. \quad (9.2)$$

Ultimately, we are interested in comparing the number of affine regions induced by the networks (P_0, N_0) and (P_1, N_1) . By Corollary 7.1.8, this comes down to comparing the number of vertices in the upper convex hulls of $P_0 \diamond N_0$ and $P_1 \diamond N_1$.

A careful analysis will reveal that the network is expected to transition from the exponential to the subexponential complexity regime. To reach this conclusion, we first develop some theoretical tools to better understand upper convex hulls arising from sums of circles.

9.1 Summing Circles

Start by enumerating the points in N_0 as

$$N_0 = \{\mathbf{x}_0, \dots, \mathbf{x}_n\},$$

where

$$\mathbf{x}_j := \left(\cos \left(\frac{\pi}{2} \frac{2j}{2n+1} \right), \sin \left(\frac{\pi}{2} \frac{2j}{2n+1} \right) \right) \in \mathbb{R}^2 \quad (9.3)$$

for $j = 0, \dots, n$.

The goal of this section is to prove the following proposition:

Proposition 9.1.1. *It holds that*

$$\mathcal{U}^*(N_0 \diamond N_0) = 2N_0.$$

To do so, we need to develop some machinery.

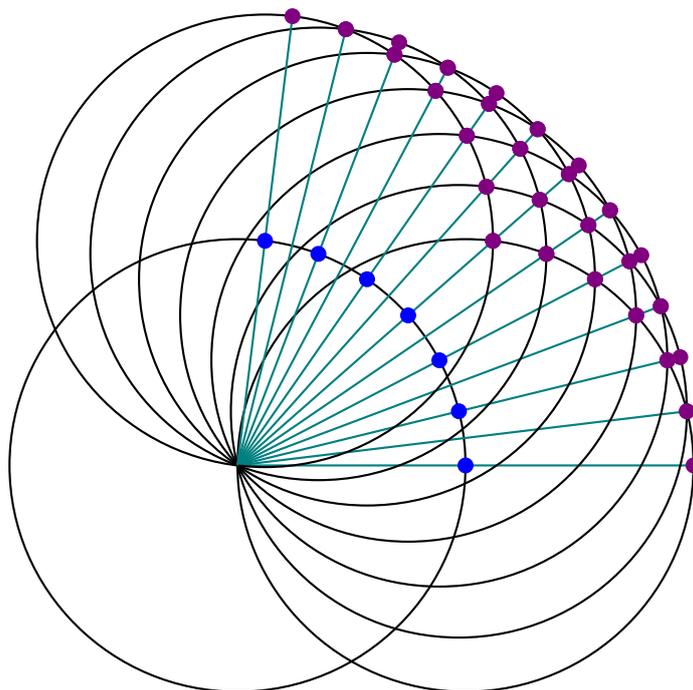


Figure 9.2: Structure underlying the Minkowski sum $N_0 \diamond N_0$, which consists of a copy of N_0 (violet) centered at n for every $n \in N_0$ (blue). The resulting set $N_0 \diamond N_0$ can be structured using lines through the origin (light blue, see Lemma 9.1.2).

We start with the following lemma (see Figure 9.2 for a visualization):

Lemma 9.1.2. *(Line Lemma)*

- i) For any two pairs $0 \leq i, j \leq n$ and $0 \leq h, k \leq n$ of indices (not necessarily distinct, i.e., it may be that $i = j$), the points $\mathbf{x}_i + \mathbf{x}_j$ and $\mathbf{x}_h + \mathbf{x}_k$ lie on a line through the origin if and only if $i + j = h + k$.
- ii) The line $0 - (\mathbf{x}_i + \mathbf{x}_j)$ corresponding to the index-pair (i, j) encloses the angle $\frac{\pi}{2} \frac{i+j}{2n+1}$ with the horizontal axis.
- iii) On that line, the point furthest away from the origin is given by $2\mathbf{x}_i$ if $i + j$ is even and $\mathbf{x}_i + \mathbf{x}_{i+1}$ if $i + j$ is odd.

Proof. We start with i). The points lie on a line through the origin if and only if

$$\frac{(\mathbf{x}_i + \mathbf{x}_j)_1}{(\mathbf{x}_h + \mathbf{x}_k)_1} = \frac{(\mathbf{x}_i + \mathbf{x}_j)_2}{(\mathbf{x}_h + \mathbf{x}_k)_2}.$$

Inserting the points, this is the case if and only if

$$\frac{\cos\left(\frac{\pi}{2} \frac{2i}{2n+1}\right) + \cos\left(\frac{\pi}{2} \frac{2j}{2n+1}\right)}{\cos\left(\frac{\pi}{2} \frac{2h}{2n+1}\right) + \cos\left(\frac{\pi}{2} \frac{2k}{2n+1}\right)} = \frac{\sin\left(\frac{\pi}{2} \frac{2i}{2n+1}\right) + \sin\left(\frac{\pi}{2} \frac{2j}{2n+1}\right)}{\sin\left(\frac{\pi}{2} \frac{2h}{2n+1}\right) + \sin\left(\frac{\pi}{2} \frac{2k}{2n+1}\right)}.$$

By Lemma A.0.2, this equation can be re-written as

$$\frac{\cos\left(\frac{\pi}{2} \frac{2i+2j}{2(2n+1)}\right) \cos\left(\frac{\pi}{2} \frac{2i-2j}{2(2n+1)}\right)}{\cos\left(\frac{\pi}{2} \frac{2h+2k}{2(2n+1)}\right) \cos\left(\frac{\pi}{2} \frac{2h-2k}{2(2n+1)}\right)} = \frac{\sin\left(\frac{\pi}{2} \frac{2i+2j}{2(2n+1)}\right) \cos\left(\frac{\pi}{2} \frac{2i-2j}{2(2n+1)}\right)}{\sin\left(\frac{\pi}{2} \frac{2h+2k}{2(2n+1)}\right) \cos\left(\frac{\pi}{2} \frac{2h-2k}{2(2n+1)}\right)},$$

which is the same as

$$\tan\left(\frac{\pi}{2} \frac{2i+2j}{2(2n+1)}\right) = \tan\left(\frac{\pi}{2} \frac{2h+2k}{2(2n+1)}\right).$$

Since $0 \leq i+j, h+k \leq 2n$, this, in turn, is equivalent to

$$i+j = h+k.$$

The first statement follows.

For the second statement, note that the angle Φ enclosed by the horizontal axis and the line defined by the index-pair (i, j) satisfies

$$\tan \Phi = \frac{(\mathbf{x}_i + \mathbf{x}_j)_2}{(\mathbf{x}_i + \mathbf{x}_j)_1}.$$

Analogously to the above calculations, one can compute

$$\tan \Phi = \tan\left(\frac{\pi}{2} \frac{2i+2j}{2(2n+1)}\right)$$

and thus

$$\Phi = \frac{\pi}{2} \frac{i+j}{2n+1}.$$

For the third claim, we compute (as above)

$$\|\mathbf{x}_i + \mathbf{x}_j\|_2^2 = 4 \cos\left(\frac{\pi}{2} \frac{i-j}{2n+1}\right).$$

If $i+j$ is even, then this term is maximized by $i=j$. If $i+j$ is odd, then this term is minimized by $|i-j|=1$. This concludes the proof. \square

The following simple Lemmas will help to prove Proposition 9.1.1

Lemma 9.1.3. It holds that $N_0 \diamond N_0 \subseteq 2D^2 := \{\mathbf{x} \in \mathbb{R}^2 \mid \|\mathbf{x}\|_2 \leq 2\}$. Furthermore, for any $0 \leq i \leq n$, it is true that $2\mathbf{x}_i \in 2S^1 := \{\mathbf{x} \in \mathbb{R}^2 \mid \|\mathbf{x}\|_2 = 2\}$.

Proof. Trivial. □

Definition 9.1.4. For any $0 \leq i \leq n - 1$, we define A_i to be the area enclosed by $2S^1$ and the line connecting $2\mathbf{x}_i$ and $2\mathbf{x}_{i+1}$ (see Figure 9.3).

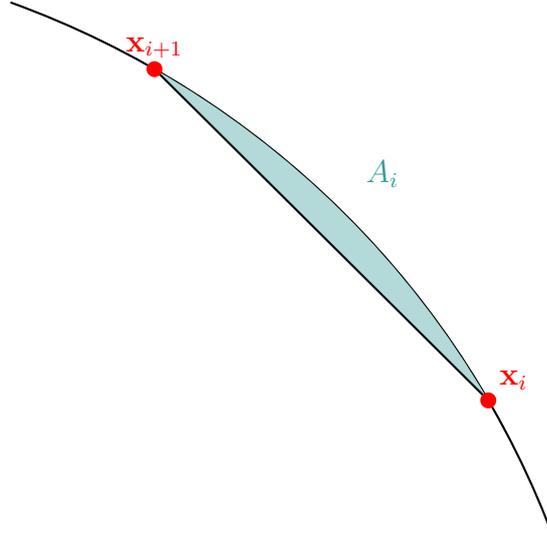


Figure 9.3: Example of the area A_i (shaded region).

As it turns out, the interior of A_i contains no points from $N_0 \diamond N_0$:

Lemma 9.1.5. For any $0 \leq i \leq n$,

$$A_i \cap (N_0 \diamond N_0) = \{2\mathbf{x}_i, 2\mathbf{x}_{i+1}, \mathbf{x}_i + \mathbf{x}_{i+1}\}.$$

Furthermore, $\mathbf{x}_i + \mathbf{x}_{i+1}$ lies on the line connecting $2\mathbf{x}_i$ and $2\mathbf{x}_{i+1}$.

Proof. The second claim is clear, as $\mathbf{x}_i + \mathbf{x}_{i+1}$ is a convex combination of $2\mathbf{x}_i$ and $2\mathbf{x}_{i+1}$, namely

$$\mathbf{x}_i + \mathbf{x}_{i+1} = \frac{1}{2}2\mathbf{x}_i + \frac{1}{2}2\mathbf{x}_{i+1}. \quad (9.4)$$

For the first claim, note that, by Lemma 9.1.2.i), the points in $N_0 \diamond N_0$ are organized on lines through the origin.

By Lemma 9.1.2.ii), there is exactly one such line running between the lines defined by $2\mathbf{x}_i$ and $2\mathbf{x}_{i+1}$, which runs through $\mathbf{x}_i + \mathbf{x}_{i+1}$. This line lies exactly between the two neighboring lines. By Lemma 9.1.2.iii), the top point of that line is $\mathbf{x}_i + \mathbf{x}_{i+1}$. It hence suffices to show that $\mathbf{x}_i + \mathbf{x}_{i+1}$ lies on the boundary of A_i , which is clear by Equation (9.4). This concludes the proof. □

We can now finally provide the...

... proof of Proposition 9.1.1. By Lemma 9.1.2, the points in $N_0 \diamond N_0$ are organized in lines, alternatingly ending in $2\mathbf{x}_i$ and $\mathbf{x}_i + \mathbf{x}_{i+1}$ (see also Figure 9.4).

By Lemma 9.1.5, the point $\mathbf{x}_i + \mathbf{x}_{i+1}$ is a convex combination of $2\mathbf{x}_i$ and $2\mathbf{x}_{i+1}$, s.t. $\mathcal{U}^*(N_0 \diamond N_0) \subseteq 2N_0$.

By Lemma 9.1.3, the set $2N_0$ is contained in $2S^1$, which implies that no point in $2N_0$ can be written as a convex combination of the other points. This shows that $2N_0 \subseteq \mathcal{U}^*(N_0 \diamond N_0)$ and concludes the proof. \square

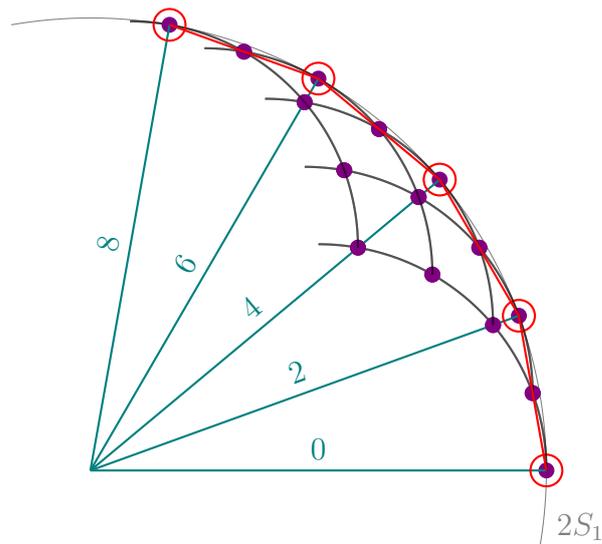


Figure 9.4: Example of $N_0 \diamond N_0$ with $n = 4$. Index-sum are included for every second line. Upper convex hull is indicated in red.

9.2 Summing Rotated Circles

In the previous section, we computed the upper convex hull of $N_0 \diamond N_0$. In this section, we study the upper convex hull of $N_0 \diamond P_0$.

Start by enumerating the points in P_0 as

$$P_0 = \{\mathbf{y}_0, \dots, \mathbf{y}_n\}$$

where

$$\mathbf{y}_j := \left(\cos \left(\frac{\pi 2j + 1}{2 2n + 1} \right), \sin \left(\frac{\pi 2j + 1}{2 2n + 1} \right) \right) \in \mathbb{R}^2, \quad (9.5)$$

for $j = 0, \dots, n$.

The goal of this section is to prove the following proposition:

Proposition 9.2.1. *It holds that*

$$\mathcal{U}^*(P_0 \diamond N_0) = \bigcup_{i=1}^n \{\mathbf{x}_i + \mathbf{y}_{i-1}, \mathbf{x}_i + \mathbf{y}_i\} \cup \{\mathbf{x}_0 + \mathbf{y}_0\}.$$

Analogously to Lemma 9.1.2 for $N_0 \diamond N_0$, the following lemma structures the points in $P_0 \diamond N_0$.

Lemma 9.2.2 (Line Lemma 2). *i) For any two pairs $0 \leq i, j \leq n$ and $0 \leq h, k \leq n$ of indices (not necessarily distinct, i.e., it may be that $i = j$), the points $\mathbf{x}_i + \mathbf{y}_j$ and $\mathbf{x}_h + \mathbf{y}_k$ lie on a line through the origin if and only if $i + j = h + k$.*

ii) The line $0 - (\mathbf{x}_i + \mathbf{y}_j)$ corresponding to the index-pair (i, j) encloses the angle $\frac{\pi}{2} \frac{i+j+1/2}{2n+1}$ with the horizontal axis.

iii) On that line, the point furthest away from the origin is given by $\mathbf{x}_i + \mathbf{y}_i$ if $i + j$ is even and $\mathbf{y}_{i-1} + \mathbf{x}_i$ if $i + j$ is odd.

Proof. We proceed analogously to the proof of Lemma 9.1.2, starting with *i)*. The two points lie on a line through the origin if and only if

$$\frac{(\mathbf{x}_i + \mathbf{y}_j)_1}{(\mathbf{x}_h + \mathbf{y}_k)_1} = \frac{(\mathbf{x}_i + \mathbf{y}_j)_2}{(\mathbf{x}_h + \mathbf{y}_k)_2}.$$

Inserting the points, this is the case if and only if

$$\frac{\cos\left(\frac{\pi}{2} \frac{2i}{2n+1}\right) + \cos\left(\frac{\pi}{2} \frac{2j+1}{2n+1}\right)}{\cos\left(\frac{\pi}{2} \frac{2h}{2n+1}\right) + \cos\left(\frac{\pi}{2} \frac{2k+1}{2n+1}\right)} = \frac{\sin\left(\frac{\pi}{2} \frac{2i}{2n+1}\right) + \sin\left(\frac{\pi}{2} \frac{2j+1}{2n+1}\right)}{\sin\left(\frac{\pi}{2} \frac{2h}{2n+1}\right) + \sin\left(\frac{\pi}{2} \frac{2k+1}{2n+1}\right)}.$$

By Lemma A.0.2, this can be re-written as

$$\frac{\cos\left(\frac{\pi}{2} \frac{2i+(2j+1)}{2(2n+1)}\right) \cos\left(\frac{\pi}{2} \frac{2i-(2j+1)}{2(2n+1)}\right)}{\cos\left(\frac{\pi}{2} \frac{2h+(2k+1)}{2(2n+1)}\right) \cos\left(\frac{\pi}{2} \frac{2h-(2k+1)}{2(2n+1)}\right)} = \frac{\sin\left(\frac{\pi}{2} \frac{2i+(2j+1)}{2(2n+1)}\right) \cos\left(\frac{\pi}{2} \frac{2i-(2j+1)}{2(2n+1)}\right)}{\sin\left(\frac{\pi}{2} \frac{2h+(2k+1)}{2(2n+1)}\right) \cos\left(\frac{\pi}{2} \frac{2h-(2k+1)}{2(2n+1)}\right)},$$

which is equivalent to

$$\tan\left(\frac{\pi}{2} \frac{2h + 2k + 1}{2(2n + 1)}\right) = \tan\left(\frac{\pi}{2} \frac{2i + 2j + 1}{2(2n + 1)}\right).$$

Since $0 \leq i + j, h + k \leq 2n$, this, in turn, is equivalent to

$$h + k = i + j.$$

The first statement follows.

For the second statement, note that the angle Φ enclosed by the horizontal axis and the line defined by the index-pair (i, j) satisfies

$$\tan \Phi = \frac{(\mathbf{x}_i + \mathbf{y}_j)_2}{(\mathbf{x}_i + \mathbf{y}_j)_1}.$$

Analogously to the above calculations, one can reformulate the right hand side to read

$$\tan \Phi = \tan\left(\frac{\pi}{2} \frac{2i + 2j + 1}{2(2n + 1)}\right)$$

and thus

$$\Phi = \frac{\pi i + j + 1/2}{2 \cdot 2n + 1}.$$

For the third statement, note that

$$\|\mathbf{x}_i + \mathbf{y}_j\|_2^2 = 4 \cos\left(\frac{\pi i - j - 1/2}{2 \cdot 2n + 1}\right).$$

If $i + j$ is even, this term is maximized for $i = j$. If $i + j$ is odd, then it is maximized by $i - j = 1$. \square

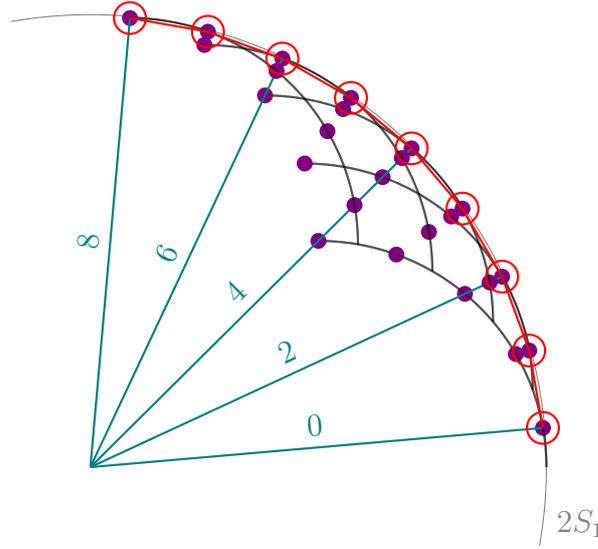


Figure 9.5: Example for $P_0 \diamond N_0$ with $n = 4$. Index-sum associated with the lines are indicated for every second line and so is the upper convex hull (red). Compared to the picture of $N_0 \diamond N_0$ (see Figure 9.4), the points are shifted counter-clockwise along the circular arcs.

The proof of Proposition 9.2.1 requires one more lemma:

Lemma 9.2.3. *Let $\mathbf{v}, \mathbf{w} \in \mathbb{R}^2$ be two vectors. Then the path $\mathbf{v} + \mathbf{w}$ obtained by appending \mathbf{w} to \mathbf{v} takes a left turn if and only if*

$$\mathbf{v} \times \mathbf{w} := \mathbf{v}_1 \mathbf{w}_2 - \mathbf{v}_2 \mathbf{w}_1 > 0.$$

Proof. Embed \mathbf{v} and \mathbf{w} into the first two dimension of \mathbb{R}^3 like

$$\tilde{\mathbf{v}} := \begin{pmatrix} \mathbf{v} \\ 0 \end{pmatrix}, \quad \tilde{\mathbf{w}} := \begin{pmatrix} \mathbf{w} \\ 0 \end{pmatrix}.$$

The cross-product

$$\tilde{\mathbf{v}} \times \tilde{\mathbf{w}} = \begin{pmatrix} 0 \\ 0 \\ \mathbf{v}_1 \mathbf{w}_2 - \mathbf{v}_2 \mathbf{w}_1 \end{pmatrix}$$

equals the *signed* area of the parallelogram spanned by $\tilde{\mathbf{v}}$ and $\tilde{\mathbf{w}}$. By our choice of orientation, this shows that P does indeed take a left turn if and only if $\mathbf{v} \times \mathbf{w} > 0$. \square

Proof of Proposition 9.2.1. “ \subseteq ”: We start by arguing that

$$\mathcal{U}^*(P_0 \diamond N_0) \subseteq \bigcup_{i=1}^n \{\mathbf{x}_i + \mathbf{y}_{i-1}, \mathbf{x}_i + \mathbf{y}_i\} \cup \{\mathbf{x}_0 + \mathbf{y}_0\}.$$

But this follows from the organization of the points in $P_0 \diamond N_0$ in lines through the origin and the fact that $\bigcup_{i=1}^n \{\mathbf{x}_i + \mathbf{y}_{i-1}, \mathbf{x}_i + \mathbf{y}_i\} \cup \{\mathbf{x}_0 + \mathbf{y}_0\}$ contains exactly the upper ends of these lines (see Lemma 9.2.2).

“ \supseteq ”:

 We show the remaining inclusion using an argument similar to the one in Andrew’s monotone chain algorithm¹. In particular, we show that the path

$$P_i: \mathbf{x}_i + \mathbf{y}_i \rightarrow \mathbf{x}_{i+1} + \mathbf{y}_i \rightarrow \mathbf{x}_{i+1} + \mathbf{y}_{i+1}$$

takes a left turn for every $i = 0, \dots, n-1$. We claim that this would show the remaining conclusion.

Indeed, since $\mathbf{x}_0 + \mathbf{y}_0$ is the rightmost point in $P_0 \diamond N_0$, it has to be contained in $\mathcal{U}^*(P_0 \diamond N_0)$. It then follows from Andrew’s argument that $\{\mathbf{x}_i + \mathbf{y}_i, \mathbf{x}_{i+1} + \mathbf{y}_i, \mathbf{x}_{i+1} + \mathbf{y}_{i+1}\} \subseteq \mathcal{U}^*(P_0, N_0)$ for $i = 0$, and thus for any i by induction.

We now show that P_i indeed takes a left turn. To do so, define the two sections in P_i as

$$\begin{aligned} \mathbf{v}_i &:= \mathbf{x}_{i+1} + \mathbf{y}_i - (\mathbf{x}_i + \mathbf{y}_i) = \mathbf{x}_{i+1} - \mathbf{x}_i \\ \mathbf{w}_i &:= \mathbf{x}_{i+1} + \mathbf{y}_{i+1} - (\mathbf{x}_i + \mathbf{y}_i) = \mathbf{y}_{i+1} - \mathbf{y}_i. \end{aligned}$$

By Lemma 9.2.3, the path P_i takes a left turn if and only if $\mathbf{v}_i \times \mathbf{w}_i = \mathbf{v}_{i1} \mathbf{w}_{i2} - \mathbf{v}_{i2} \mathbf{w}_{i1} < 0$.

Inserting the points and using Lemma A.0.3, we compute the first summand:

$$\begin{aligned} \mathbf{v}_{i1} \mathbf{w}_{i2} &= (\mathbf{x}_{i+1} - \mathbf{x}_i)_1 (\mathbf{y}_{i+1} - \mathbf{y}_i)_2 \\ &= \left[\cos \left(\frac{\pi}{2} \frac{2(i+1)}{2n+1} \right) - \cos \left(\frac{\pi}{2} \frac{2i}{2n+1} \right) \right] \left[\sin \left(\frac{\pi}{2} \frac{2(i+1)+1}{2n+1} \right) - \sin \left(\frac{\pi}{2} \frac{2i+1}{2n+1} \right) \right] \\ &\stackrel{\text{A.0.3}}{=} -4 \sin \left(\frac{\pi}{2} \frac{2(2i+1)}{2(2n+1)} \right) \sin \left(\frac{\pi}{2} \frac{2}{2(2n+1)} \right) \cos \left(\frac{\pi}{2} \frac{2(2(i+1))}{2(2n+1)} \right) \sin \left(\frac{\pi}{2} \frac{2}{2(2n+1)} \right) \end{aligned}$$

from which we conclude that

$$\mathbf{v}_{i1} \mathbf{w}_{i2} = -4 \sin \left(\frac{\pi}{2} \frac{2i+1}{2n+1} \right) \sin^2 \left(\frac{\pi}{2} \frac{1}{2n+1} \right) \cos \left(\frac{\pi}{2} \frac{2(i+1)}{2n+1} \right). \quad (9.6)$$

For the remaining summand, we proceed analogously:

$$\begin{aligned} \mathbf{v}_{i2} \mathbf{w}_{i1} &= (\mathbf{x}_{i+1} - \mathbf{x}_i)_2 (\mathbf{y}_{i+1} - \mathbf{y}_i)_1 \\ &= \left[\sin \left(\frac{\pi}{2} \frac{2(i+1)}{2n+1} \right) - \sin \left(\frac{\pi}{2} \frac{2i}{2n+1} \right) \right] \left[\cos \left(\frac{\pi}{2} \frac{2(i+1)+1}{2n+1} \right) - \cos \left(\frac{\pi}{2} \frac{2i+1}{2n+1} \right) \right] \\ &\stackrel{\text{A.0.3}}{=} -4 \cos \left(\frac{\pi}{2} \frac{2(2i+1)}{2(2n+1)} \right) \sin \left(\frac{\pi}{2} \frac{2}{2(2n+1)} \right) \sin \left(\frac{\pi}{2} \frac{2(2(i+1))}{2(2n+1)} \right) \sin \left(\frac{\pi}{2} \frac{2}{2(2n+1)} \right) \end{aligned}$$

¹Andrew’s monotone chain algorithm for computing convex hulls: https://en.wikibooks.org/wiki/Algorithm_Implementation/Geometry/Convex_hull/Monotone_chain

from it follows that

$$\mathbf{v}_{i2}\mathbf{w}_{i1} = -4 \cos\left(\frac{\pi}{2} \frac{2i+1}{2n+1}\right) \sin^2\left(\frac{\pi}{2} \frac{1}{2n+1}\right) \sin\left(\frac{\pi}{2} \frac{2(i+1)}{2n+1}\right). \quad (9.7)$$

Putting together Equation (9.6) and Equation (9.7) shows that $\mathbf{v}_i \times \mathbf{w}_i > 0$ if and only if

$$-4 \sin\left(\frac{\pi}{2} \frac{2i+1}{2n+1}\right) \cos\left(\frac{\pi}{2} \frac{2(i+1)}{2n+1}\right) + 4 \cos\left(\frac{\pi}{2} \frac{2i+1}{2n+1}\right) \sin\left(\frac{\pi}{2} \frac{2(i+1)}{2n+1}\right) > 0,$$

which can be reformulated to

$$\tan\left(\frac{\pi}{2} \frac{2(i+1)}{2n+1}\right) > \tan\left(\frac{\pi}{2} \frac{2i+1}{2n+1}\right).$$

This is true since \tan is strictly increasing on $(0, \pi)$. □

9.3 Putting It All Together

In this section, we combine the results from Section 9.1 and Section 9.2 to understand how the complexity of the network (P_0, N_0) changes when appending the random ReLU layer.

Differentiate two cases.

Case 1: $w > 0$

In this case, by Equations (9.1) and (9.2),

$$\begin{aligned} N_1 &= wN_0 \\ P_1 &= (wP_0 \boxplus b) \cup wN_0 \end{aligned}$$

and thus

$$P_1 \diamond N_1 = (wN_0 \diamond wN_0) \cup (wN_0 \diamond wP_0 \boxplus b). \quad (9.8)$$

We are interested in the upper convex hull of $P_1 \diamond N_1$. On a high level, w scales the circle on which N_0 and P_0 are arranged, while b shifts $wN_0 \diamond wP_0$ along the vertical axis. The rest of this subsection makes these considerations more precise. For ease of notation, we introduce the random-variable

$$S := |\mathcal{U}^*(P_1 \diamond N_1)| \in \mathbb{N}_+, \quad (9.9)$$

which, by Theorem 7.0.5, counts the d -cells. By Corollary 7.1.8, it furthermore is an upper bound for the number of affine regions induced by $\mathcal{Q}(P_1) - \mathcal{Q}(N_1)$.

We start with a special case, which is closely related to the previous two sections:

Proposition 9.3.1. *It holds that*

$$\mathcal{U}^*((N_0 \diamond N_0) \cup (P_0 \diamond N_0)) = \mathcal{U}^*(N_0 \diamond N_0) \cup \mathcal{U}^*(P_0 \diamond N_0).$$

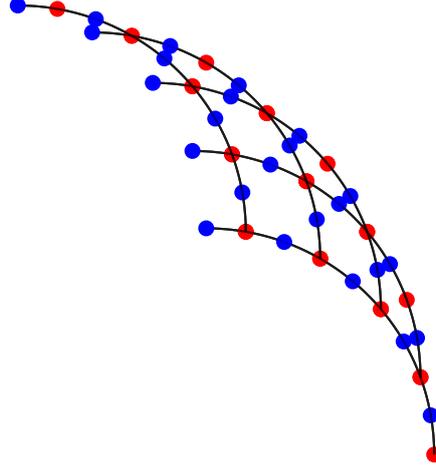


Figure 9.6: Example for $(N_0 \diamond N_0) \cup (P_0 \diamond N_0)$ with $n = 4$ ($N_0 \diamond N_0$ is red and $P_0 \diamond N_0$ is blue).

Proof. The set $(N_0 \diamond N_0) \cup (P_0 \diamond N_0)$ inherits a line-structure from $N_0 \diamond N_0$ and $P_0 \diamond N_0$ (see Lemma 9.1.2 and Lemma 9.2.2). This proposition claims that the upper convex hull of their union is the same as the union of the top points from all lines (see also Figure 9.6).

“ \subseteq ”: This is always true.

“ \supseteq ”: By Proposition 9.1.1 and Lemma 9.1.2, the set $\mathcal{U}^*(N_0 \diamond N_0)$ consists of points on lines through the origin enclosing the angle

$$\varphi_i = \frac{\pi}{2} \frac{2i}{2n+1}, \quad i = 0, \dots, n,$$

with the horizontal axis.

Furthermore, by Proposition 9.2.1 and Lemma 9.2.2, the set $\mathcal{U}^*(P_0 \diamond N_0)$ consists of points on lines through the origin enclosing the angles

$$\psi_i = \frac{\pi}{2} \frac{2i-1}{2n+1} \quad \text{or} \quad \alpha_i = \frac{\pi}{2} \frac{2i+1}{2n+1}, \quad i = 1, \dots, n$$

along with $\mathbf{x}_0 + \mathbf{y}_0$.

Consequently, the set $\mathcal{U}^*(N_0 \diamond N_0) \cup \mathcal{U}^*(P_0 \diamond N_0)$ consists of triples $(\mathbf{x}_i + \mathbf{y}_{i-1}, 2\mathbf{x}_i, \mathbf{x}_i + \mathbf{y}_i)$, $i = 1, \dots, n$, where one point from $\mathcal{U}^*(N_0 \diamond N_0)$ is enclosed by two points from $\mathcal{U}^*(P_0 \diamond N_0)$, as well as the points $\mathbf{x}_0 + \mathbf{y}_0$ and $2\mathbf{x}_0$ (the two rightmost points in Figure 9.6). Using this structure, we will show the remaining inclusion.

Note that $2\mathbf{x}_0 \in \mathcal{U}^*((N_0 \diamond N_0) \cup (P_0 \diamond N_0))$ since it is the right-most point.

Next, the point $\mathbf{x}_0 + \mathbf{y}_0$ is contained in $\mathcal{U}^*((N_0 \diamond N_0) \cup (P_0 \diamond N_0))$ if and only if the path $P: 2\mathbf{x}_0 \rightarrow \mathbf{x}_0 + \mathbf{y}_0 \rightarrow \mathbf{x}_1 + \mathbf{y}_0$ takes a left turn. To see that this is the case, define

$$\begin{aligned} \mathbf{v} &:= \mathbf{x}_0 + \mathbf{y}_0 - 2\mathbf{x}_0 = \mathbf{y}_0 - \mathbf{x}_0 \\ \mathbf{w} &:= \mathbf{x}_1 + \mathbf{y}_0 - (\mathbf{x}_0 + \mathbf{y}_0) = \mathbf{x}_1 - \mathbf{x}_0. \end{aligned}$$

Then

$$\begin{aligned}
\mathbf{v}_1 \mathbf{w}_2 &= (\mathbf{y}_0 - \mathbf{x}_0)_1 (\mathbf{x}_1 - \mathbf{x}_0)_2 \\
&= \left(\cos \left(\frac{\pi}{2} \frac{1}{2n+1} \right) - \cos(0) \right) \left(\sin \left(\frac{\pi}{2} \frac{2}{2n+1} \right) - \sin(0) \right) \\
&= \left(\cos \left(\frac{\pi}{2} \frac{1}{2n+1} \right) - 1 \right) \sin \left(\frac{\pi}{2} \frac{2}{2n+1} \right)
\end{aligned}$$

and, analogously,

$$\begin{aligned}
\mathbf{v}_2 \mathbf{w}_1 &= (\mathbf{y}_0 - \mathbf{x}_0)_2 (\mathbf{x}_1 - \mathbf{x}_0)_1 \\
&= \left(\sin \left(\frac{\pi}{2} \frac{1}{2n+1} \right) - \sin(0) \right) \left(\cos \left(\frac{\pi}{2} \frac{2}{2n+1} \right) - \cos(0) \right) \\
&= \sin \left(\frac{\pi}{2} \frac{1}{2n+1} \right) \left(\cos \left(\frac{\pi}{2} \frac{2}{2n+1} \right) - 1 \right).
\end{aligned}$$

Thus, the path makes a left turn if and only if

$$\left(\cos \left(\frac{\pi}{2} \frac{1}{2n+1} \right) - 1 \right) \sin \left(\frac{\pi}{2} \frac{2}{2n+1} \right) > \sin \left(\frac{\pi}{2} \frac{1}{2n+1} \right) \left(\cos \left(\frac{\pi}{2} \frac{2}{2n+1} \right) - 1 \right),$$

which is equivalent to

$$\begin{aligned}
\sin \left(\frac{\pi}{2} \frac{1}{2n+1} \right) &= \sin \left(\frac{\pi}{2} \frac{2}{2n+1} \right) \cos \left(\frac{\pi}{2} \frac{1}{2n+1} \right) - \sin \left(\frac{\pi}{2} \frac{2}{2n+1} \right) \cos \left(\frac{\pi}{2} \frac{1}{2n+1} \right) \\
&> \sin \left(\frac{\pi}{2} \frac{2}{2n+1} \right) - \sin \left(\frac{\pi}{2} \frac{1}{2n+1} \right).
\end{aligned}$$

The claim then follows from the fact that

$$2 \sin \left(\frac{\pi}{2} \frac{1}{2n+1} \right) > \sin \left(\frac{\pi}{2} \frac{2}{2n+1} \right).$$

Finally, we are left to show that the triples are contained in the upper convex hull. This is the case if and only if the paths $P_i := w\mathbf{x}_i + w\mathbf{y}_{i-1} \rightarrow 2w\mathbf{x}_i \rightarrow w\mathbf{x}_i + w\mathbf{y}_i$ make a left turn (see Figure 9.7).

Define

$$\begin{aligned}
\mathbf{v}_i &:= 2w\mathbf{x}_i - (w\mathbf{x}_i + w\mathbf{y}_{i-1}) = w\mathbf{x}_i - w\mathbf{y}_{i-1} \\
\mathbf{w}_i &:= w\mathbf{x}_i + w\mathbf{y}_i - 2w\mathbf{x}_i = w\mathbf{y}_i - w\mathbf{x}_i.
\end{aligned}$$

Then

$$\begin{aligned}
\mathbf{v}_{i1} \mathbf{w}_{i2} &= (w\mathbf{x}_i - w\mathbf{y}_{i-1})_1 (w\mathbf{y}_i - w\mathbf{x}_i)_2 \\
&= w^2 \left[\cos \left(\frac{\pi}{2} \frac{2i}{2n+1} \right) - \cos \left(\frac{\pi}{2} \frac{2i-1}{2n+1} \right) \right] \left[\sin \left(\frac{\pi}{2} \frac{2i+1}{2n+1} \right) - \sin \left(\frac{\pi}{2} \frac{2i}{2n+1} \right) \right] \\
&= -4w^2 \sin \left(\frac{\pi}{2} \frac{4i-1}{4n+2} \right) \sin \left(\frac{\pi}{2} \frac{1}{4n+2} \right) \cos \left(\frac{\pi}{2} \frac{4i+1}{4n+2} \right) \sin \left(\frac{\pi}{2} \frac{1}{4n+2} \right)
\end{aligned}$$

and

$$\begin{aligned}
\mathbf{v}_{i2} \mathbf{w}_{i1} &= (w\mathbf{x}_i - w\mathbf{y}_{i-1})_2 (w\mathbf{y}_i - w\mathbf{x}_i)_1 \\
&= w^2 \left[\sin\left(\frac{\pi}{2} \frac{2i}{2n+1}\right) - \sin\left(\frac{\pi}{2} \frac{2i-1}{2n+1}\right) \right] \left[\cos\left(\frac{\pi}{2} \frac{2i+1}{2n+1}\right) - \cos\left(\frac{\pi}{2} \frac{2i}{2n+1}\right) \right] \\
&= -4w^2 \cos\left(\frac{\pi}{2} \frac{4i-1}{4n+2}\right) \sin\left(\frac{\pi}{2} \frac{1}{4n+2}\right) \sin\left(\frac{\pi}{2} \frac{4i+1}{4n+2}\right) \sin\left(\frac{\pi}{2} \frac{1}{4n+2}\right),
\end{aligned}$$

which implies that

$$\begin{aligned}
\mathbf{v}_i \times \mathbf{w}_i &= -4w^2 \sin^2\left(\frac{\pi}{2} \frac{1}{4n+2}\right) \left(\sin\left(\frac{\pi}{2} \frac{4i-1}{4n+2}\right) \cos\left(\frac{\pi}{2} \frac{4i+1}{4n+2}\right) - \right. \\
&\quad \left. - \cos\left(\frac{\pi}{2} \frac{4i-1}{4n+2}\right) \sin\left(\frac{\pi}{2} \frac{4i+1}{4n+2}\right) \right) \\
&\stackrel{A.0.4}{=} 4w^2 \sin^2\left(\frac{\pi}{2} \frac{1}{4n+2}\right) \sin\left(\frac{\pi}{2} \frac{2}{4n+2}\right)
\end{aligned}$$

and thus $\mathbf{v}_i \times \mathbf{w}_i > 0$. This concludes the proof. \square

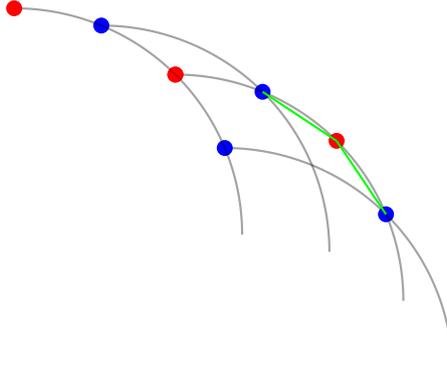


Figure 9.7: Path $P_i := w\mathbf{x}_i + w\mathbf{y}_{i-1} + \boxplus b \rightarrow 2w\mathbf{x}_i \rightarrow w\mathbf{x}_i + w\mathbf{y}_i \boxplus b$ (counter-clockwise, green. Here with $b = 0$). If it makes a left turn, then $2w\mathbf{x}_i$ is contained in the upper convex hull, given $w > 0$.

The following theorem generalizes Proposition 9.3.1 for arbitrary $w > 0, b > 0$:

Theorem 9.3.2. *Assume $w, b > 0$. Then*

$$\mathcal{U}^*(P_1 \boxplus N_1) = \mathcal{U}^*(wN_0 \boxplus wP_0 \boxplus b) \cup \{2w\mathbf{x}_0\} \cup \bigcup_{\substack{1 \leq i \leq n: \\ u(i)w > b}} \{2w\mathbf{x}_i\}, \quad (9.10)$$

where

$$u(i) := \frac{2 \sin^2\left(\frac{\pi}{2} \frac{1}{2(2n+1)}\right)}{\sin\left(\frac{\pi}{2} \frac{4i}{2(2n+1)}\right)}. \quad (9.11)$$

Proof. Proposition 9.3.1 describes the special case where $w = 1$ and $b = 0$. The proposition generalizes to any $w > 0$ for $b = 0$:

$$\mathcal{U}^* \left((wN_0 \diamond wN_0) \cup (wN_0 \diamond wP_0) \right) = \mathcal{U}^* \left((wN_0 \diamond wN_0) \right) \cup \mathcal{U}^* \left((wN_0 \diamond wP_0) \right). \quad (9.12)$$

In this case, it is clear that $u^{(i)}w > b = 0$ for all $i = 1, \dots, n$ and hence Equation (9.10) holds. The interesting part is the influence of $b > 0$.

So assume $b > 0$. Then, by Equation (9.8), the set $P_1 \diamond N_1$ consists of $wN_0 \diamond wN_0$ along with $wN_0 \diamond wP_0$ shifted upwards by b . This means that, upon considering $b > 0$, the term $\mathcal{U}^* \left((wN_0 \diamond wP_0) \boxplus b \right)$ from Equation (9.12) will always be contained in $\mathcal{U}^*(P_1 + N_1)$:

$$\mathcal{U}^*(P_1 \diamond N_1) \supseteq \mathcal{U}^*(wN_0 \diamond wP_0 \boxplus b).$$

Furthermore

$$\mathcal{U}^*(P_1 \diamond N_1) \supseteq \{2w\mathbf{x}_0\},$$

since $2w\mathbf{x}_0$ is the left-most point in $P_1 \cup N_1$.

It remains to study which part of the upper convex hull of $wN_0 \diamond wN_0$ is also contained in the upper convex hull of $P_1 \diamond N_1$. Similar to the proof of Proposition 9.3.1, we study the conditions under which the path $P_i := w\mathbf{x}_i + w\mathbf{y}_{i-1} \boxplus b \rightarrow 2w\mathbf{x}_i \rightarrow w\mathbf{x}_i + w\mathbf{y}_i \boxplus b$ makes a left turn (see Figure 9.7). By an Andrew-type argument, this is the case if and only if $2w\mathbf{x}_i \in P_1 \diamond N_1$.

Define

$$\begin{aligned} \mathbf{v}_i &:= 2w\mathbf{x}_i - (w\mathbf{x}_i + w\mathbf{y}_{i-1} \boxplus b) = w\mathbf{x}_i - w\mathbf{y}_{i-1} \boxplus -b \\ \mathbf{w}_i &:= w\mathbf{x}_i + w\mathbf{y}_i \boxplus b - 2w\mathbf{x}_i = w\mathbf{y}_i - w\mathbf{x}_i \boxplus b. \end{aligned}$$

Then

$$\begin{aligned} \mathbf{v}_{i1}\mathbf{w}_{i2} &= (w\mathbf{x}_i - w\mathbf{y}_{i-1} \boxplus -b)_1(w\mathbf{y}_i - w\mathbf{x}_i \boxplus b)_2 \\ &= w^2 \left[\cos \left(\frac{\pi}{2} \frac{2i}{2n+1} \right) - \cos \left(\frac{\pi}{2} \frac{2i-1}{2n+1} \right) \right] \left[\sin \left(\frac{\pi}{2} \frac{2i+1}{2n+1} \right) - \sin \left(\frac{\pi}{2} \frac{2i}{2n+1} \right) + \frac{b}{w} \right] \\ &= -2w \sin \left(\frac{\pi}{2} \frac{4i-1}{4n+2} \right) \sin \left(\frac{\pi}{2} \frac{1}{4n+2} \right) \left[2w \cos \left(\frac{\pi}{2} \frac{4i+1}{4n+2} \right) \sin \left(\frac{\pi}{2} \frac{1}{4n+2} \right) + b \right] \end{aligned}$$

and

$$\begin{aligned} \mathbf{v}_{i2}\mathbf{w}_{i1} &= (w\mathbf{x}_i - w\mathbf{y}_{i-1} \boxplus -b)_2(w\mathbf{y}_i - w\mathbf{x}_i \boxplus b)_1 \\ &= w^2 \left[\sin \left(\frac{\pi}{2} \frac{2i}{2n+1} \right) - \sin \left(\frac{\pi}{2} \frac{2i-1}{2n+1} \right) - \frac{b}{w} \right] \left[\cos \left(\frac{\pi}{2} \frac{2i+1}{2n+1} \right) - \cos \left(\frac{\pi}{2} \frac{2i}{2n+1} \right) \right] \\ &= \left[2w \cos \left(\frac{\pi}{2} \frac{4i-1}{4n+2} \right) \sin \left(\frac{\pi}{2} \frac{1}{4n+2} \right) - b \right] \left[-2w \sin \left(\frac{\pi}{2} \frac{4i+1}{4n+2} \right) \sin \left(\frac{\pi}{2} \frac{1}{4n+2} \right) \right], \end{aligned}$$

which implies that

$$\begin{aligned}
\mathbf{v}_i \times \mathbf{w}_i &= -4w^2 \sin^2 \left(\frac{\pi}{2} \frac{1}{4n+2} \right) \left(\sin \left(\frac{\pi}{2} \frac{4i-1}{4n+2} \right) \cos \left(\frac{\pi}{2} \frac{4i+1}{4n+2} \right) - \right. \\
&\quad \left. - \cos \left(\frac{\pi}{2} \frac{4i-1}{4n+2} \right) \sin \left(\frac{\pi}{2} \frac{4i+1}{4n+2} \right) \right) \\
&\quad - 2wb \sin \left(\frac{\pi}{2} \frac{1}{4n+2} \right) \left(\sin \left(\frac{\pi}{2} \frac{4i-1}{4n+2} \right) + \sin \left(\frac{\pi}{2} \frac{4i+1}{4n+2} \right) \right) \\
&\stackrel{A.0.4}{=} 4w^2 \sin^2 \left(\frac{\pi}{2} \frac{1}{4n+2} \right) \sin \left(\frac{\pi}{2} \frac{2}{4n+2} \right) - \\
&\quad - 2bw \sin \left(\frac{\pi}{2} \frac{1}{4n+2} \right) \left(\sin \left(\frac{\pi}{2} \frac{4i-1}{4n+2} \right) + \sin \left(\frac{\pi}{2} \frac{4i+1}{4n+2} \right) \right) \\
&\stackrel{A.0.2}{=} 4w^2 \sin^2 \left(\frac{\pi}{2} \frac{1}{4n+2} \right) \sin \left(\frac{\pi}{2} \frac{2}{4n+2} \right) - \\
&\quad - 2bw \sin \left(\frac{\pi}{2} \frac{1}{4n+2} \right) 2 \sin \left(\frac{\pi}{2} \frac{8i}{2(4n+2)} \right) \cos \left(\frac{\pi}{2} \frac{2}{2(4n+2)} \right) \\
&= 4w^2 \sin^2 \left(\frac{\pi}{2} \frac{1}{4n+2} \right) \sin \left(\frac{\pi}{2} \frac{2}{4n+2} \right) - \\
&\quad - 4bw \sin \left(\frac{\pi}{2} \frac{1}{4n+2} \right) \cos \left(\frac{\pi}{2} \frac{1}{4n+2} \right) \sin \left(\frac{\pi}{2} \frac{4i}{4n+2} \right),
\end{aligned}$$

and thus $\mathbf{v}_i \times \mathbf{w}_i > 0$ if and only if

$$\begin{aligned}
b &< \frac{\sin^2 \left(\frac{\pi}{2} \frac{1}{4n+2} \right) \sin \left(\frac{\pi}{2} \frac{2}{4n+2} \right)}{\sin \left(\frac{\pi}{2} \frac{1}{4n+2} \right) \cos \left(\frac{\pi}{2} \frac{1}{4n+2} \right) \sin \left(\frac{\pi}{2} \frac{4i}{4n+2} \right)} w \\
&= \frac{2 \sin^2 \left(\frac{\pi}{2} \frac{1}{2(2n+1)} \right)}{\sin \left(\frac{\pi}{2} \frac{4i}{2(2n+1)} \right)} w = \mathbf{u}(i)w.
\end{aligned}$$

□

The previous theorem tells us how the upper convex hull looks if $w > 0, b > 0$. The following following theorem handles the case $w > 0, b < 0$:

Theorem 9.3.3. *Assume $w > 0, b < 0$. Then*

$$\mathcal{U}^*(P_1 \diamond N_1) = \mathcal{U}^*(wN_0 \diamond wN_0 \boxplus b) \cup \{w(\mathbf{x}_n + \mathbf{y}_n)\} \cup \bigcup_{\substack{0 \leq i \leq n-1: \\ \mathbf{l}(i)w < b}} \{w(\mathbf{x}_i + \mathbf{y}_i) \boxplus b, w(\mathbf{x}_i + \mathbf{y}_{i+1}) \boxplus b\}, \tag{9.13}$$

where

$$\mathbf{l}(i) := - \frac{2 \sin^2 \left(\frac{\pi}{2} \frac{1}{2(2n+1)} \right)}{\sin \left(\frac{\pi}{2} \frac{2(2i+1)}{2(2n+1)} \right)} \tag{9.14}$$

Proof. The proof of this theorem is very similar to the proof of Theorem 9.3.2, but instead of upwards, the set $wN_0 \diamond wP_0$ is shifted downwards by b . Analogously, one can see that

$$\mathcal{U}^*(P_1 \diamond N_1) \supseteq \mathcal{U}^*(wN_0 \diamond wN_0) \cup \{w\mathbf{x}_0 + w\mathbf{y}_0\}.$$

It remains to study which points in $wN_0 \diamond wP_0 \boxplus b$ are also part of the upper convex hull of $P_1 \diamond N_1$.

We start with the observation that, for each $i = 1, \dots, n-1$, the line l_1^i from $2w\mathbf{x}_i$ to $2w\mathbf{x}_{i+1}$ and the line l_2^i from $w\mathbf{x}_i + w\mathbf{y}_i$ to $w\mathbf{x}_{i+1} + w\mathbf{y}_i$ are parallel (Figure 9.8). Indeed, this follows from the fact that $w\mathbf{x}_i + w\mathbf{y}_i - (w\mathbf{x}_{i+1} + w\mathbf{y}_i) = w\mathbf{x}_i - w\mathbf{x}_{i+1}$.

Thus, the points $w\mathbf{x}_i + w\mathbf{y}_i \boxplus b$ and $w\mathbf{x}_{i+1} + w\mathbf{y}_i \boxplus b$ lie in the upper convex hull of $P_1 \diamond N_1$ as long as l_2^i lies above l_1^i . We employ an Andrew-type argument to derive a condition under which that is guaranteed.

Let $P_i: 2w\mathbf{x}_i \rightarrow w\mathbf{x}_i + w\mathbf{y}_i \boxplus b \rightarrow w\mathbf{x}_{i+1} + w\mathbf{y}_i \boxplus b$ and define

$$\begin{aligned} \mathbf{v}_i &:= w\mathbf{x}_i + w\mathbf{y}_i \boxplus b - 2w\mathbf{x}_i = w\mathbf{y}_i - w\mathbf{x}_i \boxplus b \\ \mathbf{w}_i &:= w\mathbf{x}_{i+1} + w\mathbf{y}_i \boxplus b - (w\mathbf{x}_i + w\mathbf{y}_i \boxplus b) = w\mathbf{x}_{i+1} - w\mathbf{x}_i. \end{aligned}$$

The goal is figuring out when this path makes a left turn.

We compute

$$\begin{aligned} \mathbf{v}_{i1}\mathbf{w}_{i2} &= (w\mathbf{y}_i - w\mathbf{x}_i \boxplus b)_1(w\mathbf{x}_{i+1} - w\mathbf{x}_i)_2 \\ &= w^2 \left[\cos\left(\frac{\pi}{2} \frac{2i+1}{2n+1}\right) - \cos\left(\frac{\pi}{2} \frac{2i}{2n+1}\right) \right] \left[\sin\left(\frac{\pi}{2} \frac{2(i+1)}{2n+1}\right) - \sin\left(\frac{\pi}{2} \frac{2i}{2n+1}\right) \right] \\ &= -4w^2 \sin\left(\frac{\pi}{2} \frac{4i+1}{4n+2}\right) \sin\left(\frac{\pi}{2} \frac{1}{4n+2}\right) \cos\left(\frac{\pi}{2} \frac{2(2i+1)}{4n+2}\right) \sin\left(\frac{\pi}{2} \frac{2}{4n+2}\right) \end{aligned}$$

and

$$\begin{aligned} \mathbf{v}_{i2}\mathbf{w}_{i1} &= (w\mathbf{y}_i - w\mathbf{x}_i \boxplus b)_2(w\mathbf{x}_{i+1} - w\mathbf{x}_i)_1 \\ &= w^2 \left[\sin\left(\frac{\pi}{2} \frac{2i+1}{2n+1}\right) - \sin\left(\frac{\pi}{2} \frac{2i}{2n+1}\right) + \frac{b}{w} \right] \left[\cos\left(\frac{\pi}{2} \frac{2i+2}{2n+1}\right) - \cos\left(\frac{\pi}{2} \frac{2i}{2n+1}\right) \right] \\ &= \left[2w \cos\left(\frac{\pi}{2} \frac{4i+1}{4n+2}\right) \sin\left(\frac{\pi}{2} \frac{1}{4n+2}\right) + b \right] \left[-2w \sin\left(\frac{\pi}{2} \frac{4i+2}{4n+2}\right) \sin\left(\frac{\pi}{2} \frac{2}{4n+2}\right) \right] \\ &= -4w^2 \cos\left(\frac{\pi}{2} \frac{4i+1}{4n+2}\right) \sin\left(\frac{\pi}{2} \frac{1}{4n+2}\right) \sin\left(\frac{\pi}{2} \frac{4i+2}{4n+2}\right) \sin\left(\frac{\pi}{2} \frac{2}{4n+2}\right) - \\ &\quad - 2wb \sin\left(\frac{\pi}{2} \frac{4i+2}{4n+2}\right) \sin\left(\frac{\pi}{2} \frac{2}{4n+2}\right), \end{aligned}$$

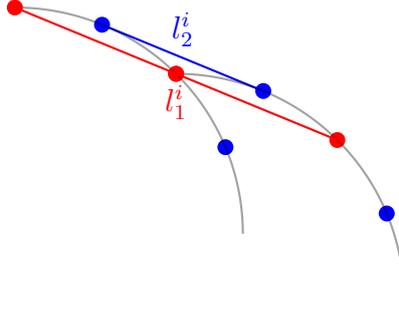


Figure 9.8: The two parallel lines from $2w\mathbf{x}_i$ to $2w\mathbf{x}_{i+1}$ (l_1^i , red) and $w\mathbf{y}_i + w\mathbf{y}_i$ to $w\mathbf{x}_{i+1} + w\mathbf{y}_i$ (l_2^i , blue).

which implies that

$$\begin{aligned} \mathbf{v}_i \times \mathbf{w}_i &= -4w^2 \sin\left(\frac{\pi}{2} \frac{1}{4n+2}\right) \sin\left(\frac{\pi}{2} \frac{2}{4n+2}\right) \\ &\cdot \left[\sin\left(\frac{\pi}{2} \frac{4i+1}{4n+2}\right) \cos\left(\frac{\pi}{2} \frac{4i+2}{4n+2}\right) - \cos\left(\frac{\pi}{2} \frac{4i+1}{4n+2}\right) \sin\left(\frac{\pi}{2} \frac{4i+2}{4n+2}\right) \right] + \\ &+ 2bw \sin\left(\frac{\pi}{2} \frac{4i+2}{4n+2}\right) \sin\left(\frac{\pi}{2} \frac{2}{4n+2}\right) \\ &= 4w^2 \sin^2\left(\frac{\pi}{2} \frac{1}{4n+2}\right) \sin\left(\frac{\pi}{2} \frac{2}{4n+2}\right) + 2bw \sin\left(\frac{\pi}{2} \frac{4i+2}{4n+2}\right) \sin\left(\frac{\pi}{2} \frac{2}{4n+2}\right), \end{aligned}$$

and thus $\mathbf{v}_i \times \mathbf{v}_i > 0$ if and only if

$$b > -\frac{2 \sin^2\left(\frac{\pi}{2} \frac{1}{2(2n+1)}\right)}{\sin\left(\frac{\pi}{2} \frac{2(2i+1)}{2(2n+1)}\right)} w = \mathfrak{l}(i)w.$$

This concludes the proof. \square

Remark 9.3.4. Theorem 9.3.3 provides an example for a non-empty set of paths $\mathfrak{P}(P_1 \diamond N_1)$. In particular, if b is small enough, there will be two adjacent points $p_i + n_i = 2\mathbf{x}_i$, $p_{i+1} + n_{i+1} = 2\mathbf{x}_{i+1}$ in the upper convex hull of $P_1 \diamond N_1$ which clearly satisfy $p_i - n_i = p_{i+1} - n_{i+1}$. Consequently, the path $(p_i + n_i, p_{i+1} + n_{i+1})$ is contained in $\mathfrak{P}(P_1 \diamond N_1)$. In particular, this shows that $|\mathcal{U}^*(P_1 \diamond N_1)|$ is a strict upper bound for the number of affine regions defined by $\mathcal{Q}(P_1) - \mathcal{Q}(N_1)$ (see Corollary 7.1.8).

Theorem 9.3.2 and Theorem 9.3.3 enable us to count the points in the upper convex hull:

Lemma 9.3.5. For a fixed n , $u(\cdot, n)$ is positive and decreasing while $\mathfrak{l}(\cdot, n)$ is negative and increasing.

Theorem 9.3.6. The conditional random variable $S \mid w > 0, b > 0$ (with S defined as in Equation (9.9)) takes the values

$$S \mid w > 0, b > 0 \in \{2n + 2 + i \mid 0 \leq i \leq n\} \quad (9.15)$$

with probability

$$\mathbb{P}(S = 2n + 2 + i \mid w > 0, b > 0) = \begin{cases} 1 - \frac{2}{\pi} \tan^{-1}(|\mathbf{u}(1)|), & i = 0 \\ \frac{2}{\pi} (\tan^{-1}(|\mathbf{u}(i)|) - \tan^{-1}(|\mathbf{u}(i+1)|)), & 0 < i < n \\ \frac{2}{\pi} \tan^{-1}(|\mathbf{u}(n)|), & i = n. \end{cases} \quad (9.16)$$

Proof. By Theorem 9.3.2, $\mathcal{U}^*(P_1 \diamond N_1)$ consists of $\mathcal{U}^*(wN_0 \diamond wP_0 \boxplus b) \cup \{2w\mathbf{x}_0\}$, regardless of the exact value of w and b , together with a set depending on w and b . By Proposition 9.2.1,

$$|\mathcal{U}^*(wN_0 \diamond wP_0 \boxplus b)| = 2n + 1$$

(indeed, w just scales the points and b shifts them along the y -axis. In particular, the total size of this upper convex hull is independent of the values of $w > 0$ and $b > 0$). This shows Equation (9.15).

We now compute the conditional distribution. Let $1 < i < n$. It follows from Theorem 9.3.2 and Lemma 9.3.5 that

$$\begin{aligned} \mathbb{P}(S \geq 2n + 2 + i \mid w > 0, b > 0) &= \mathbb{P}(b < \mathbf{u}(i)w \mid w > 0, b > 0) \\ &= \frac{\mathbb{P}(b < \mathbf{u}(i)w, w > 0, b > 0)}{\mathbb{P}(w > 0, b > 0)} \\ &= 4\mathbb{P}(0 < b < \mathbf{u}(i)w) \end{aligned}$$

and thus, with $s := 2n + 2 + i$,

$$\begin{aligned} \mathbb{P}(S = s \mid w > 0, b > 0) &= \mathbb{P}(S \geq s \mid w > 0, b > 0) - \mathbb{P}(S \geq s + 1 \mid w > 0, b > 0) \\ &= 4(\mathbb{P}(0 < b < \mathbf{u}(i)w) - \mathbb{P}(0 < b < \mathbf{u}(i+1, n)w)) \\ &= \frac{4}{2\pi} (\tan^{-1}(|\mathbf{u}(i)|) - \tan^{-1}(|\mathbf{u}(i+1)|)), \end{aligned}$$

where in the last step we used Lemma A.0.1. Furthermore, for $i = n$, we compute

$$\begin{aligned} \mathbb{P}(S = 3n + 2 \mid w > 0, b > 0) &= \mathbb{P}(S \geq 3n + 2 \mid w > 0, b > 0) \\ &= 4\mathbb{P}(0 < b < \mathbf{u}(n)w) \\ &= \frac{4}{2\pi} \tan^{-1}(|\mathbf{u}(n)|) \end{aligned}$$

and

$$\begin{aligned} \mathbb{P}(S = 2n + 2 \mid w > 0, b > 0) &= 1 - \mathbb{P}(S \geq 2n + 2 + 1 \mid w > 0, b > 0) \\ &= 1 - 4\mathbb{P}(0 < b < \mathbf{u}(1)) \\ &= 1 - \frac{4}{2\pi} \tan^{-1}(|\mathbf{u}(1)|). \end{aligned}$$

□

Theorem 9.3.7. *The random variable $S \mid w > 0, b < 0$ takes the values*

$$S \mid w > 0, b < 0 \in \{n + 2 + 2i \mid 0 \leq i \leq n\} \quad (9.17)$$

with probability

$$\mathbb{P}(S = n + 2 + 2i \mid w > 0, b < 0) = \begin{cases} 1 - \frac{2}{\pi} \tan^{-1} (|\mathfrak{l}(0)|), & i = 0 \\ \frac{2}{\pi} (\tan^{-1} (|\mathfrak{l}(i-1)|) - \tan^{-1} (|\mathfrak{l}(i)|)), & 0 < i < n \\ \frac{2}{\pi} \tan^{-1} (|\mathfrak{l}(n-1)|), & i = n. \end{cases} \quad (9.18)$$

Proof. It follows from Theorem 9.3.3 that $\mathcal{U}^*(P_1 \diamond N_1)$ consists of $\mathcal{U}^*(wN_0 \diamond wN_0) \cup \{w\mathbf{x}_0 + w\mathbf{y}_0\}$, independent of the values of b and w , along with a set of points depending on w and b . Since

$$|\mathcal{U}^*(wN_0 \diamond wN_0)| = n + 1$$

by Proposition 9.1.1, Equation (9.17) follows.

We now compute the conditional distribution. Let $0 < i < n$. It follows from Theorem 9.3.3 and Lemma 9.3.5 that, for $s := n + 2 + 2i$,

$$\begin{aligned} \mathbb{P}(S = s \mid w > 0, b < 0) &= \mathbb{P}(S \geq s \mid w > 0, b < 0) - \mathbb{P}(S \geq s + 2 \mid w > 0, b < 0) \\ &= 4\mathbb{P}(0 > b > \mathfrak{l}(i-1, n)w) - 4\mathbb{P}(0 > b > \mathfrak{l}(i)w). \end{aligned}$$

Furthermore,

$$\mathbb{P}(S = 3n + 2 \mid w > 0, b < 0) = \mathbb{P}(S \geq 3n + 2 \mid w > 0, b < 0)$$

and

$$\begin{aligned} \mathbb{P}(S = n + 2 \mid w > 0, b < 0) &= 1 - \mathbb{P}(S \geq n + 2 + 2 \mid w > 0, b < 0) \\ &= 1 - 4\mathbb{P}(S \geq n + 2 + 2 \mid w > 0, b < 0). \end{aligned}$$

The claim then follows from Lemma A.0.1. □

This concludes our study of the case $w < 0$. The next subsection deals with $w < 0$.

Case 2: $w < 0$

In this case, by Equations (9.1) and (9.2),

$$\begin{aligned} N_1 &= w^- P_0 \\ P_1 &= (w^- N_0 \boxplus b) \cup N_1 \end{aligned}$$

and thus

$$P_1 \diamond N_1 = (w^-(P_0 \diamond P_0)) \cup (w^-(N_0 \boxplus b + P_0)). \quad (9.19)$$

Like before, when studying the case $w > 0$, we start by establishing the special case $b = 0$:

Proposition 9.3.8. *Assume $w < 0$. Then it holds that*

$$\mathcal{U}^*(w^-(P_0 \diamond P_0) \cup w^-(P_0 \diamond N_0)) = \mathcal{U}^*(w^-(P_0 \diamond P_0)) \cup \mathcal{U}^*(w^-(P_0 \diamond N_0)).$$

Proof. Analogous to the proof of Proposition 9.3.1. □

The following theorem generalizes this result to arbitrary $b > 0$:

Theorem 9.3.9. *Assume that $w < 0$ and $b > 0$. Then*

$$\mathcal{U}^*(P_1 \diamond N_1) = \mathcal{U}^*(w^-N_0 \diamond w^-P_0 \boxplus b) \cup \{2w^-y_0\} \cup \bigcup_{\substack{0 \leq i \leq n-1 \\ -l(n,i)w^- > b}} \{2w^-y_i\}. \quad (9.20)$$

Proof. Proposition 9.3.8 explains the special case when $b = 0$. The case $b > 0$ arises from the case $b = 0$ by shifting the points in $w^-N_0 \diamond w^-P_0$ upward along the vertical axis by b . Similarly to the proof of Theorem 9.3.2, depending on the magnitude of b , the point $2w^-y_i$ will also be contained in the upper convex hull. This is the case if and only if the path $P_i: w^-(x_i + y_i) \boxplus b \rightarrow 2w^-y_i \rightarrow w^-(x_{i+1} + y_i) \boxplus b$ takes a left turn.

Define

$$\begin{aligned} \mathbf{v}_i &:= 2w^-y_i - (w^-(x_i + y_i) \boxplus b) = w^-(y_i - x_i) \boxplus -b \\ \mathbf{w}_i &:= w^-(x_{i+1} + y_i) \boxplus b - 2w^-y_i = w^-(x_{i+1} - y_i) \boxplus b. \end{aligned}$$

Then

$$\begin{aligned} \mathbf{v}_{i1}\mathbf{w}_{i2} &= (w^-y_i - w^-x_i \boxplus -b)_1(w^-x_{i+1} - w^-y_i \boxplus b)_2 \\ &= w^{-2} \left[\cos\left(\frac{\pi}{2} \frac{2i+1}{2n+1}\right) - \cos\left(\frac{\pi}{2} \frac{2i}{2n+1}\right) \right] \left[\sin\left(\frac{\pi}{2} \frac{2i+2}{2n+1}\right) - \sin\left(\frac{\pi}{2} \frac{2i+1}{2n+1}\right) + \frac{b}{w^-} \right] \\ &= -2w^- \sin\left(\frac{\pi}{2} \frac{4i+1}{4n+2}\right) \sin\left(\frac{\pi}{2} \frac{1}{4n+2}\right) \left[2w^- \cos\left(\frac{\pi}{2} \frac{4i+3}{4n+2}\right) \sin\left(\frac{\pi}{2} \frac{1}{4n+2}\right) + b \right] \end{aligned}$$

and

$$\begin{aligned} \mathbf{v}_{i2}\mathbf{w}_{i1} &= (w^-y_i - w^-x_i \boxplus -b)_2(w^-x_{i+1} - w^-y_i \boxplus b)_1 \\ &= w^{-2} \left[\sin\left(\frac{\pi}{2} \frac{2i+1}{2n+1}\right) - \sin\left(\frac{\pi}{2} \frac{2i}{2n+1}\right) - \frac{b}{w^-} \right] \left[\cos\left(\frac{\pi}{2} \frac{2i+2}{2n+1}\right) - \cos\left(\frac{\pi}{2} \frac{2i+1}{2n+1}\right) \right] \\ &= \left[2w^- \cos\left(\frac{\pi}{2} \frac{4i+1}{4n+2}\right) \sin\left(\frac{\pi}{2} \frac{1}{4n+2}\right) - b \right] \left[-2w^- \sin\left(\frac{\pi}{2} \frac{4i+3}{4n+2}\right) \sin\left(\frac{\pi}{2} \frac{1}{4n+2}\right) \right] \end{aligned}$$

which implies

$$\begin{aligned}
\mathbf{v}_i \times \mathbf{w}_i &= -4w^{-2} \sin^2 \left(\frac{\pi}{2} \frac{1}{2(2n+1)} \right) \left(\sin \left(\frac{\pi}{2} \frac{4i+1}{2(2n+1)} \right) \cos \left(\frac{\pi}{2} \frac{4i+3}{2(2n+1)} \right) - \right. \\
&\quad \left. - \cos \left(\frac{\pi}{2} \frac{4i+1}{2(2n+1)} \right) \sin \left(\frac{\pi}{2} \frac{4i+3}{2(2n+1)} \right) \right) \\
&\quad - 2w^{-b} \sin \left(\frac{\pi}{2} \frac{1}{2(2n+1)} \right) \left(\sin \left(\frac{\pi}{2} \frac{4i+1}{2(2n+1)} \right) + \sin \left(\frac{\pi}{2} \frac{4i+3}{2(2n+1)} \right) \right) \\
&\stackrel{A.0.4}{=} 4w^{-2} \sin^2 \left(\frac{\pi}{2} \frac{1}{2(2n+1)} \right) \sin \left(\frac{\pi}{2} \frac{2}{2(2n+1)} \right) - \\
&\quad - 2bw^{-b} \sin \left(\frac{\pi}{2} \frac{1}{2(2n+1)} \right) \left(\sin \left(\frac{\pi}{2} \frac{4i+1}{2(2n+1)} \right) + \sin \left(\frac{\pi}{2} \frac{4i+3}{2(2n+1)} \right) \right) \\
&= 4w^{-2} \sin^2 \left(\frac{\pi}{2} \frac{1}{2(2n+1)} \right) \sin \left(\frac{\pi}{2} \frac{2}{2(2n+1)} \right) - \\
&\quad - 2bw^{-b} \sin \left(\frac{\pi}{2} \frac{1}{2(2n+1)} \right) 2 \sin \left(\frac{\pi}{2} \frac{8i+4}{2(2(2n+1))} \right) \cos \left(\frac{\pi}{2} \frac{2}{2(2(2n+1))} \right) \\
&= 4w^{-2} \sin^2 \left(\frac{\pi}{2} \frac{1}{2(2n+1)} \right) \sin \left(\frac{\pi}{2} \frac{2}{2(2n+1)} \right) - \\
&\quad - 4bw^{-b} \sin \left(\frac{\pi}{2} \frac{1}{2(2n+1)} \right) \cos \left(\frac{\pi}{2} \frac{1}{2(2n+1)} \right) \sin \left(\frac{\pi}{2} \frac{4i+2}{2(2n+1)} \right),
\end{aligned}$$

and thus $\mathbf{v}_{i1}\mathbf{w}_{i2} - \mathbf{v}_{i2}\mathbf{w}_{i1} > 0$ if and only if

$$\begin{aligned}
b &< \frac{\sin^2 \left(\frac{\pi}{2} \frac{1}{2(2n+1)} \right) \sin \left(\frac{\pi}{2} \frac{2}{2(2n+1)} \right)}{\sin \left(\frac{\pi}{2} \frac{1}{2(2n+1)} \right) \cos \left(\frac{\pi}{2} \frac{1}{2(2n+1)} \right) \sin \left(\frac{\pi}{2} \frac{4i+2}{2(2n+1)} \right)} w^{-} \\
&= \frac{2 \sin^2 \left(\frac{\pi}{2} \frac{1}{2(2n+1)} \right)}{\sin \left(\frac{\pi}{2} \frac{4i+2}{2(2n+1)} \right)} w^{-} = -l(i)w^{-}.
\end{aligned}$$

□

Finally, the remaining case $b < 0$:

Theorem 9.3.10. *Assume that $w < 0$ and $b < 0$. Then*

$$\mathcal{U}^*(P_1 \diamond N_1) = \mathcal{U}^*(w^{-}P_0 \diamond w^{-}P_0) \cup \{w^{-}(\mathbf{x}_0 + \mathbf{y}_0) \boxplus b\} \cup \bigcup_{\substack{1 \leq i \leq n \\ -u(n,i)w^{-} < b}} \{w^{-}(\mathbf{x}_i + \mathbf{y}_i) \boxplus b, w^{-}(\mathbf{x}_i + \mathbf{y}_{i-1})\}. \quad (9.21)$$

Proof. Proposition 9.3.8 again explains the edge case when $b = 0$. The case $b < 0$ arises from this edge case by shifting the points in $wN_0 \diamond wP_0$ downward along the vertical axis by b .

We proceed analogously to the proof of Theorem 9.3.3. Depending on the magnitude of b , the points $w^{-}(\mathbf{y}_{i-1} + \mathbf{x}_i) \boxplus b$ and $w^{-}(\mathbf{y}_i + \mathbf{x}_i) \boxplus b$ are contained in the upper convex hull.

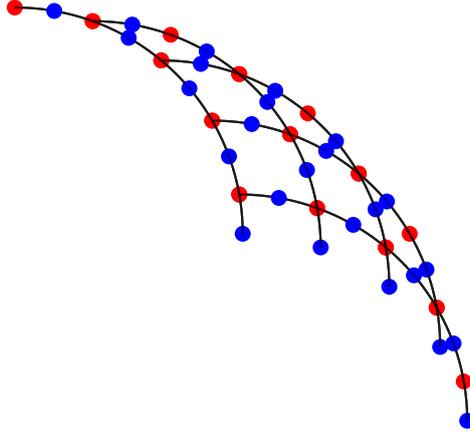


Figure 9.9: Example for $(P_0 \diamond P_0) \cup (P_0 \diamond N_0)$ with $n = 4$ ($P_0 \cup P_0$ is red and $P_0 \cup N_0$ is blue).

This is the case if and only if the path $P: 2w^-y_{i-1} \rightarrow w^-(x_i + y_{i-1}) \boxplus b \rightarrow w^-(x_i + y_i) \boxplus b$ makes a left turn.

Let

$$\begin{aligned} \mathbf{v}_i &:= w^-(x_i + y_{i-1}) \boxplus b - 2w^-y_{i-1} = w^-(x_i - y_{i-1}) \boxplus b \\ \mathbf{w}_i &:= w^-(x_i + y_i) \boxplus b - (w^-(x_i + y_{i-1}) \boxplus b) = w^-(y_i - y_{i-1}). \end{aligned}$$

Then

$$\begin{aligned} \mathbf{v}_{i1}\mathbf{w}_{i2} &= (w^-x_i - w^-y_{i-1} \boxplus b)_1(wy_i - wy_{i-1})_2 \\ &= w^{-2} \left[\cos\left(\frac{\pi}{2} \frac{2i}{2n+1}\right) - \cos\left(\frac{\pi}{2} \frac{2i-1}{2n+1}\right) \right] \left[\sin\left(\frac{\pi}{2} \frac{2i+1}{2n+1}\right) - \sin\left(\frac{\pi}{2} \frac{2i-1}{2n+1}\right) \right] \\ &= -4w^{-2} \sin\left(\frac{\pi}{2} \frac{4i-1}{4n+2}\right) \sin\left(\frac{\pi}{1} \frac{1}{4n+2}\right) \cos\left(\frac{\pi}{2} \frac{4i}{4n+2}\right) \sin\left(\frac{\pi}{2} \frac{2}{4n+2}\right) \end{aligned}$$

and

$$\begin{aligned} \mathbf{v}_{i2}\mathbf{w}_{i1} &= (w^-x_i - w^-y_{i-1} \boxplus b)_2(wy_i - wy_{i-1})_1 \\ &= w^{-2} \left[\sin\left(\frac{\pi}{2} \frac{2i}{2n+1}\right) - \sin\left(\frac{\pi}{2} \frac{2i-1}{2n+1}\right) + \frac{b}{w^-} \right] \left[\cos\left(\frac{\pi}{2} \frac{2i+1}{2n+1}\right) - \cos\left(\frac{\pi}{2} \frac{2i-1}{2n+1}\right) \right] \\ &= \left[2w^- \cos\left(\frac{\pi}{2} \frac{4i-1}{4n+2}\right) \sin\left(\frac{\pi}{2} \frac{1}{4n+2}\right) + b \right] \left[-2w^- \sin\left(\frac{\pi}{2} \frac{4i}{4n+2}\right) \sin\left(\frac{\pi}{2} \frac{2}{4n+2}\right) \right] \\ &= -4w^{-2} \cos\left(\frac{\pi}{2} \frac{4i-1}{4n+2}\right) \sin\left(\frac{\pi}{2} \frac{1}{4n+2}\right) \sin\left(\frac{\pi}{2} \frac{4i}{4n+2}\right) \sin\left(\frac{\pi}{2} \frac{2}{4n+2}\right) - \\ &\quad - 2w^-b \sin\left(\frac{\pi}{2} \frac{4i}{4n+2}\right) \sin\left(\frac{\pi}{2} \frac{2}{4n+2}\right), \end{aligned}$$

which implies that

$$\begin{aligned}
\mathbf{v}_i \times \mathbf{w}_i &= -4w^{-2} \sin\left(\frac{\pi}{2} \frac{1}{2(2n+1)}\right) \sin\left(\frac{\pi}{2} \frac{2}{2(2n+1)}\right) \\
&\quad \cdot \left(\sin\left(\frac{\pi}{2} \frac{4i-1}{4n+2}\right) \cos\left(\frac{\pi}{2} \frac{4i}{4n+2}\right) - \cos\left(\frac{\pi}{2} \frac{4i-1}{4n+2}\right) \sin\left(\frac{\pi}{2} \frac{4i}{4n+2}\right) \right) + \\
&\quad + 2bw^{-} \sin\left(\frac{\pi}{2} \frac{4i}{2(2n+1)}\right) \sin\left(\frac{\pi}{2} \frac{2}{2(2n+1)}\right) \\
&= 4w^{-2} \sin^2\left(\frac{\pi}{2} \frac{1}{2(2n+1)}\right) \sin\left(\frac{\pi}{2} \frac{2}{2(2n+1)}\right) \\
&\quad + 2bw^{-} \sin\left(\frac{\pi}{2} \frac{4i}{2(2n+1)}\right) \sin\left(\frac{\pi}{2} \frac{2}{2(2n+1)}\right)
\end{aligned}$$

and thus $\mathbf{v}_{i1}\mathbf{w}_{i2} - \mathbf{v}_{i2}\mathbf{w}_{i1} > 0$ if and only if

$$b > -\frac{2 \sin^2\left(\frac{\pi}{2} \frac{1}{2(2n+1)}\right)}{\sin\left(\frac{\pi}{2} \frac{4i}{2(2n+1)}\right)} w^{-} = -\mathbf{u}(i)w^{-}.$$

This concludes the proof. \square

Using Theorem 9.3.9 and Theorem 9.3.10, we can count the points in $\mathcal{U}^*(P_1 \diamond N_1)$, given $w < 0$:

Theorem 9.3.11. *The random variable $S \mid w < 0, b > 0$ takes the values*

$$S \mid w < 0, b > 0 \in \{2n + 2 + i \mid 0 \leq i \leq n\} \quad (9.22)$$

with probability

$$\mathbb{P}(S = 2n + 2 + i \mid w < 0, b > 0) = \begin{cases} 1 - \frac{2}{\pi} \tan^{-1}(|\mathbf{l}(0)|), & i = 0 \\ \frac{2}{\pi} (\tan^{-1}(|\mathbf{l}(i-1)|) - \tan^{-1}(|\mathbf{l}(i)|)), & 0 < i < n \\ \frac{2}{\pi} \tan^{-1}(|\mathbf{l}(n-1)|), & i = n. \end{cases} \quad (9.23)$$

Proof. Goes analogously to the proof of Theorem 9.3.6 (note that the index is shifted down by one since adding just one point ($i = 1$), this corresponds to adding $2w^{-}\mathbf{y}_0$). \square

Theorem 9.3.12. *The random variable $S \mid w < 0, b < 0$ takes the values*

$$S \mid w < 0, b < 0 \in \{n + 2 + 2i \mid 0 \leq i \leq n\} \quad (9.24)$$

with probability

$$\mathbb{P}(S = n + 2 + 2i \mid w < 0, b < 0) = \begin{cases} 1 - \frac{2}{\pi} \tan^{-1}(|\mathbf{u}(1)|), & i = 0 \\ \frac{2}{\pi} (\tan^{-1}(|\mathbf{u}(i)|) - \tan^{-1}(|\mathbf{u}(i+1)|)), & 0 < i < n \\ \frac{2}{\pi} \tan^{-1}(|\mathbf{u}(n)|), & i = n. \end{cases} \quad (9.25)$$

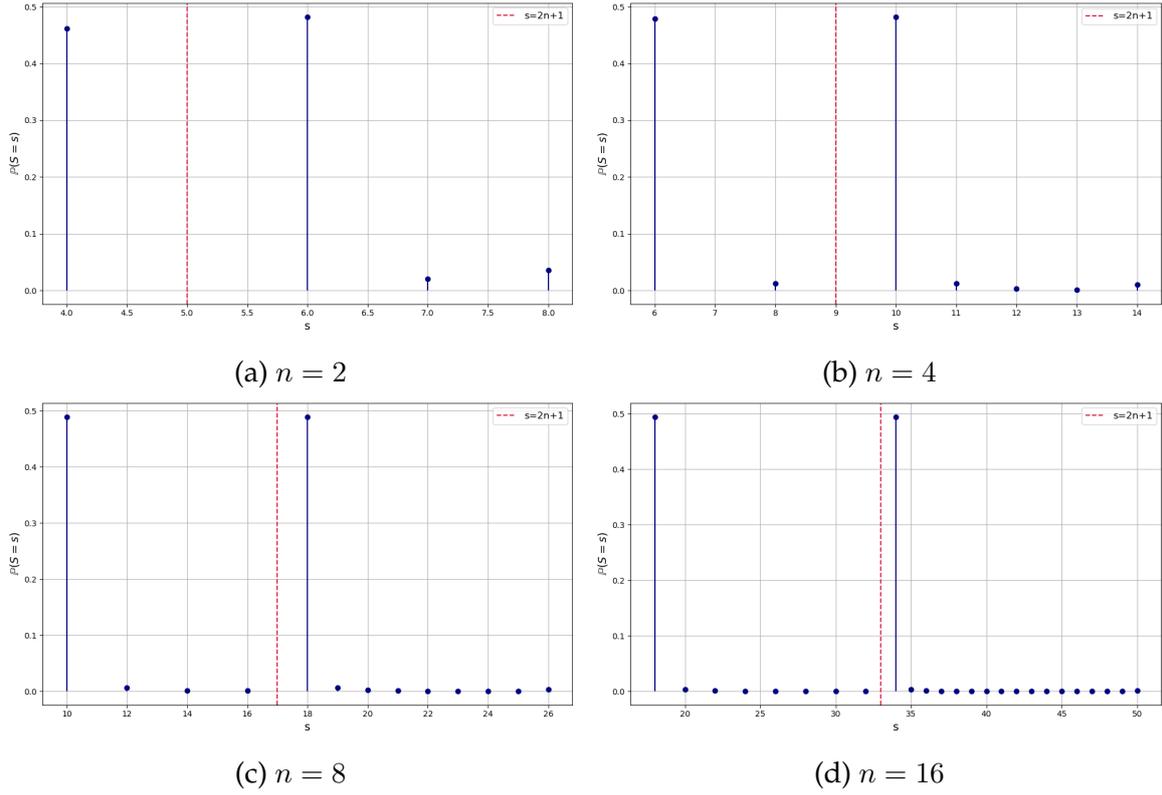


Figure 9.10: Distribution $\mathbb{P}(S)$, given in Theorem 9.3.14, for $n = 2, 4, 8, 16$.

Proof. Analogous to the proof of Theorem 9.3.7. □

The following is an interesting fact. We will not use it for the rest of our argument.

Corollary 9.3.13. *For all $0 \leq i \leq n$, it holds that*

$$\mathbb{P}(S = 2n + 2 + i | w > 0, b > 0) = \mathbb{P}(S = n + 2 + 2i | w < 0, b < 0)$$

and

$$\mathbb{P}(S = n + 2 + 2i | w > 0, b < 0) = \mathbb{P}(S = 2n + 2 + i | w < 0, b > 0).$$

Proof. Follows from Theorems 9.3.6, 9.3.7, 9.3.11 and 9.3.12. □

Bringing together all four conditional distributions of S allows writing down a closed-form expression of the unconditional distribution. In order to simplify notation, we introduce the following short-hand notations:

$$\mathbf{al} := \tan^{-1}(|\mathbf{l}|), \quad \mathbf{au} := \tan^{-1}(|\mathbf{u}|).$$

and

$$\Delta \mathbf{al}_i := \tan^{-1}(|\mathbf{l}(i)|) - \tan^{-1}(|\mathbf{l}(i+1)|), \quad \Delta \mathbf{au}_i := \tan^{-1}(|\mathbf{u}(i)|) - \tan^{-1}(|\mathbf{u}(i+1)|)$$

Using these definitions, the following theorem provides the distribution of S :

Theorem 9.3.14. *Let $n \geq 2$. Then*

$$\mathbb{P}(S = n + 2 + i) = \begin{cases} \frac{1}{2} - \frac{1}{2\pi} (\mathbf{au}(1) - \mathbf{al}(0)), & i = 0 \\ \frac{1}{2\pi} \left(\Delta \mathbf{al}_{\frac{i}{2}-1} + \Delta \mathbf{au}_{\frac{i}{2}} \right), & 0 < i < n \text{ even} \\ \frac{1}{2} - \frac{1}{2\pi} (\mathbf{au}(1) + \mathbf{al}(0) - \Delta \mathbf{al}_{\frac{n}{2}-1} - \Delta \mathbf{au}_{\frac{n}{2}}), & i = n \text{ even} \\ \frac{1}{2} - \frac{1}{2\pi} (\mathbf{au}(1) + \mathbf{al}(0)), & i = n \text{ odd} \\ \frac{1}{2\pi} \left(\Delta \mathbf{au}_{i-n} + \Delta \mathbf{al}_{i-n-1} + \Delta \mathbf{au}_{\frac{i}{2}} + \Delta \mathbf{al}_{\frac{i}{2}-1} \right), & n < i < 2n \text{ even} \\ \frac{1}{2\pi} (\Delta \mathbf{au}_{i-n} + \Delta \mathbf{al}_{i-n-1}), & n < i < 2n \text{ odd} \\ \frac{1}{2\pi} (2\mathbf{au}(n) + 2\mathbf{al}(n-1)), & i = 2n \\ 0, & o.w. \end{cases} \quad (9.26)$$

Proof. For any $s \in \mathbb{N}$,

$$\begin{aligned} \mathbb{P}(S = s) &= \mathbb{P}(S = s \mid w > 0, b > 0) \mathbb{P}(w > 0) \mathbb{P}(b > 0) + \\ &\quad + \mathbb{P}(S = s \mid w > 0, b < 0) \mathbb{P}(w > 0) \mathbb{P}(b < 0) + \\ &\quad + \mathbb{P}(S = s \mid w < 0, b > 0) \mathbb{P}(w < 0) \mathbb{P}(b > 0) + \\ &\quad + \mathbb{P}(S = s \mid w < 0, b < 0) \mathbb{P}(w < 0) \mathbb{P}(b < 0) \\ &= \frac{1}{4} \left[\mathbb{P}(S = s \mid w > 0, b > 0) + \mathbb{P}(S = s \mid w > 0, b < 0) \right. \\ &\quad \left. + \mathbb{P}(S = s \mid w < 0, b > 0) + \mathbb{P}(S = s \mid w < 0, b < 0) \right]. \end{aligned}$$

The theorem then follows from Theorems 9.3.6, 9.3.7, 9.3.11 and 9.3.12 by carefully studying all cases. \square

The distribution $\mathbb{P}(S)$ is plotted in Figure 9.10. The two main peaks correspond to $s = n + 2$ and $s = 2n + 2$. In particular, the distribution allows computing the expected complexity of the upper convex hull, which is a big step towards understanding how it behaves after one random layer. However, the resulting sum is messy and hard to comprehend. In the next section we simplify this sum.

9.4 Implications

In this section, we study the implications of the previous section for the complexity of the upper convex hull. In particular, we show that the probability of the additional layer increasing the number of d -cells approaches $1/2$ from above like $1/2 + \mathcal{O}(n^{-2})$ (Proposition 9.4.3). Furthermore, we show that the number of affine regions is expected to decrease (Corollary 9.4.14).

We start by providing the number of affine regions induced by the network $\mathcal{Q}(P_0) - \mathcal{Q}(N_0)$:

Proposition 9.4.1. *The number of affine regions induced by the network $\mathcal{Q}(P_0) - \mathcal{Q}(N_0)$ is $2n + 1$. This is also the number of d -cells in $\mathcal{T}(P_0, N_0)$.*

Proof. By Theorem 7.0.5 and Proposition 9.2.1, the number of d -cells in $\mathcal{T}(P_0, N_0)$ is

$$|\mathcal{U}^*(P_0 \diamond N_0)| = 2n + 1.$$

One can quickly confirm that $\mathfrak{B}(P_0 \diamond N_0) = \emptyset$. The proposition then follows from Corollary 7.1.8. \square

by Proposition 9.4.1, the probability $\mathbb{P}(\uparrow)$ of increasing the number of d -cells by appending the random ReLU layer is

$$\mathbb{P}(\uparrow) = \mathbb{P}(S > 2n + 1). \quad (9.27)$$

The following theorem provides a closed-form expression for this probability.

Theorem 9.4.2. *Let $n \geq 2$. Then*

$$\mathbb{P}(\uparrow) = \frac{1}{2} + \frac{2}{\pi} \left(\tan^{-1} \left(\left| \mathfrak{l} \left(\left\lceil \frac{n}{2} \right\rceil - 1 \right) \right| \right) + \tan^{-1} \left(\left| \mathfrak{u} \left(\left\lceil \frac{n}{2} \right\rceil \right) \right| \right) \right). \quad (9.28)$$

Proof. By Theorems 9.3.6, 9.3.7, 9.3.11 and 9.3.12, the probability is given by

$$\begin{aligned} \mathbb{P}(\uparrow) &= \mathbb{P}(w > 0, b > 0) + \mathbb{P}(w < 0, b > 0) + \\ &\quad + \mathbb{P}(w > 0, b < 0, S > 2n + 1) + \mathbb{P}(w < 0, b < 0, S > 2n + 1) \\ &= \frac{1}{2} + \mathbb{P}(w > 0, b < 0, S > 2n + 1) + \mathbb{P}(w < 0, b < 0, S > 2n + 1) \\ &= \frac{1}{2} + \mathbb{P}(S > 2n + 1, w > 0, b < 0) + \mathbb{P}(S > 2n + 1, w < 0, b < 0) \\ &= \frac{1}{2} + \mathbb{P} \left(0 > b > \mathfrak{l} \left(\left\lceil \frac{n}{2} \right\rceil - 1 \right) w \right) + \mathbb{P} \left(0 > b > \mathfrak{u} \left(\left\lceil \frac{n}{2} \right\rceil \right) w \right), \end{aligned}$$

where in the last step we used an argument analogous to the proof of Theorem 9.3.7 and Theorem 9.3.10. By Lemma A.0.1, this can be re-written as

$$\mathbb{P}(\uparrow) = \frac{1}{2} + \frac{2}{\pi} \left[\tan^{-1} \left(\left| \mathfrak{l} \left(\left\lceil \frac{n}{2} \right\rceil - 1 \right) \right| \right) + \tan^{-1} \left(\left| \mathfrak{u} \left(\left\lceil \frac{n}{2} \right\rceil \right) \right| \right) \right].$$

\square

Define

$$\delta(n) := \frac{2}{\pi} \left(\tan^{-1} \left(\left| \mathfrak{l} \left(\left\lceil \frac{n}{2} \right\rceil - 1 \right) \right| \right) + \tan^{-1} \left(\left| \mathfrak{u} \left(\left\lceil \frac{n}{2} \right\rceil \right) \right| \right) \right)$$

to be the non-trivial contribution to the probability of increase. The following proposition shows that $\delta(n)$ vanishes like $\mathcal{O}(n^{-2})$ for large n .

Proposition 9.4.3. *It holds that*

$$\delta(n) = \mathcal{O}(n^{-2}) \text{ as } n \rightarrow \infty.$$

Proof. It is known that

$$\sin(x) = \mathcal{O}(x) \text{ as } x \rightarrow 0,$$

which implies

$$\mathfrak{t}\left(\left\lceil\left\lfloor\frac{n}{2}\right\rfloor - 1\right\rceil\right) = \frac{2 \sin^2\left(\frac{\pi}{2} \frac{1}{2(2n+1)}\right)}{\sin\left(\frac{\pi}{2} \frac{2(2\lceil\frac{n}{2}\rceil - 1 + 1)}{2(2n+1)}\right)} = \mathcal{O}(n^{-2})$$

and

$$\mathfrak{u}\left(\left\lceil\left\lfloor\frac{n}{2}\right\rfloor\right\rceil\right) = \frac{2 \sin^2\left(\frac{\pi}{2} \frac{1}{2(2n+1)}\right)}{\sin\left(\frac{\pi}{2} \frac{4\lceil\frac{n}{2}\rceil}{2(2n+1)}\right)} = \mathcal{O}(n^{-2}).$$

The observation that $\arctan(x) = \mathcal{O}(x)$ as $x \rightarrow 0$ then concludes the proof. \square

The above argument shows that the probability of increasing the number of d -cells decreases like $\mathcal{O}(n^{-2})$ towards $1/2$. The following considerations further deepen our understanding of the complexity by computing its expectation.

We start with a number of helpful statements:

Lemma 9.4.4. *For any $m \in \mathbb{N}$, it is true that*

$$\sum_{k=1}^{m-1} \frac{1}{\sin\left(\frac{k\pi}{2m}\right)} = \frac{1}{2} \left(\sum_{k=1}^{2m-1} \frac{1}{\sin\left(\frac{k\pi}{2m}\right)} - 1 \right).$$

Proof. By symmetry, $\sin\left(\frac{k\pi}{2m}\right) = \sin\left(\frac{(2m-k)\pi}{2m}\right)$ for all $k \in \{1, \dots, m-1\}$, so that

$$\sum_{k=1}^{2m-1} \frac{1}{\sin\left(\frac{k\pi}{2m}\right)} = 2 \sum_{k=1}^{m-1} \frac{1}{\sin\left(\frac{k\pi}{2m}\right)} + \frac{1}{\sin\left(\frac{m\pi}{2m}\right)}$$

and thus

$$\sum_{k=1}^{m-1} \frac{1}{\sin\left(\frac{k\pi}{2m}\right)} = \frac{1}{2} \left(\sum_{k=1}^{2m-1} \frac{1}{\sin\left(\frac{k\pi}{2m}\right)} - 1 \right).$$

\square

We furthermore need the harmonic numbers:

Definition 9.4.5 (Harmonic Number). Given $n \in \mathbb{N}$, the *harmonic number* H_n is defined as

$$H_n := \sum_{k=1}^n \frac{1}{k}.$$

The following is a well-known result, expressing the harmonic numbers integrals:

Lemma 9.4.6. Given $n \in \mathbb{N}$, the harmonic number H_n satisfies

$$H_n = \int_0^1 \frac{1-x^n}{1-x} dx.$$

Proof. The integrand can be written as a geometric series,

$$\frac{1-x^n}{1-x} = 1 + x + \dots + x^{n-1}.$$

Integration yields

$$\int_0^1 \frac{1-x^n}{1-x} dx = \sum_{k=0}^{n-1} \frac{x^{k+1}}{k+1} \Big|_0^1 = H_n.$$

□

This result allows us to express a sum of inverse sines using the harmonic numbers:

Proposition 9.4.7 (Outlined in [31]). For any $m \in \mathbb{N}$ and $N \geq 1$, it holds that

$$\sum_{k=1}^{m-1} \csc\left(\frac{\pi k}{m}\right) = \frac{2m}{\pi} \ln\left(\frac{2m}{\pi}\right) + \frac{2m}{\pi} \gamma + \sum_{k=1}^{N-1} \frac{a_{2k}}{(2m/\pi)^{2k-1}} + C_{N,m} \frac{a_{2N}}{(2m/\pi)^{2N-1}},$$

where

$$a_{2k} := -\frac{B_{2k} 2\eta(2k)}{2k} \left(\frac{2}{\pi}\right)^{2k}$$

with the Bernoulli-numbers B_{2k} , the Dirichlet eta function η , and

$$C_{N,m} := \frac{-\sigma_{N,m} - \theta_{N,m}(2\eta(2N) - 1)}{2\eta(2N)}$$

for two bounded constants $0 < \sigma_{N,m}, \theta_{N,m} < 1$.

Proof. We elaborate on the proof outlined in [31], starting by claiming that

$$\sum_{k=1}^{m-1} \csc\left(\frac{\pi k}{m}\right) = \frac{2m}{\pi} \int_0^1 \frac{s^{m-1} - 1}{(s-1)(s^m + 1)} ds. \quad (9.29)$$

Indeed, by the well-known² representation

$$\csc(x) = \frac{1}{\pi} \int_0^\infty \frac{1}{s^2 + s} s^{x/\pi} ds$$

²<https://functions.wolfram.com/ElementaryFunctions/Csc/introductions/Csc/05/>

for any $0 < x < \pi$, so that

$$\begin{aligned} \sum_{k=1}^{m-1} \csc\left(\frac{\pi k}{n}\right) &= \frac{1}{\pi} \int_0^{\infty} \frac{1}{s^2 + s} \sum_{k=1}^{m-1} s^{k/m} ds \\ &= \frac{1}{\pi} \int_0^{\infty} \frac{\frac{s-1}{s^{1/m-1}} - 1}{s^2 + s} ds \\ &= \frac{1}{\pi} \int_0^{\infty} \frac{s - s^{1/m}}{(s^2 + s)(s^{1/m} - 1)} ds. \end{aligned}$$

After a change of variables $s = x^{-m}$ with $ds = -mx^{-m-1}dx$, this integral reads

$$\begin{aligned} \frac{1}{\pi} \int_0^{\infty} \frac{s - s^{1/m}}{(s^2 + s)(s^{1/m} - 1)} ds &= \frac{m}{\pi} \int_0^{\infty} x^{-m-1} \frac{x^{-m} - x^{-1}}{(x^{-2m} + x^{-m})(x^{-1} - 1)} dx \\ &= \frac{m}{\pi} \int_0^{\infty} \frac{1 - x^{m-1}}{(1 + x^m)(1 - x)} dx. \end{aligned}$$

Next, a quick substitution $x = t^{-1}$ confirms that

$$\int_0^1 \frac{1 - x^{m-1}}{(1 + x^m)(1 - x)} dx = \int_1^{\infty} \frac{1 - t^{m-1}}{(1 + t^m)(1 - t)} dt.$$

Equation (9.29) follows.

For the next step, we use the identity

$$\frac{s^{m-1} - 1}{(s-1)(s^m + 1)} = \frac{s^m - 1}{(s-1)(s^m + 1)} - \frac{s^{m-1}}{s^m + 1}$$

and Equation (9.29) to write

$$\begin{aligned} \sum_{k=1}^{m-1} \csc\left(\frac{\pi k}{m}\right) &= \frac{2m}{\pi} \left(\int_0^1 \frac{s^m - 1}{(s-1)(s^m + 1)} ds - \int_0^1 \frac{s^{m-1}}{s^m + 1} ds \right) \\ &= \frac{2m}{\pi} \left(\int_0^1 \frac{s^m - 1}{(s-1)(s^m + 1)} ds - \frac{\ln 2}{m} \right). \end{aligned}$$

Then

$$\frac{s^m - 1}{s-1} - \frac{s^m - 1}{s-1} \frac{s^m}{s^m + 1} = \frac{s^m - 1}{(s-1)(s^m + 1)}$$

implies

$$\sum_{k=1}^{m-1} \operatorname{csc}\left(\frac{\pi k}{m}\right) = \frac{2m}{\pi} \left(\int_0^1 \frac{s^m - 1}{s - 1} ds - \int_0^1 \frac{s^m - 1}{s - 1} \frac{s^m}{s^m + 1} ds - \frac{\ln 2}{m} \right) \quad (9.30)$$

$$= \frac{2m}{\pi} \left(H_m - \int_0^1 \frac{s^m - 1}{s - 1} \frac{s^m}{s^m + 1} ds - \frac{\ln 2}{m} \right) \quad (9.31)$$

where H_m is the m -th harmonic number (see Lemma 9.4.6).

It remains to compute the integral

$$I(m) := \int_0^1 \frac{s^m - 1}{s - 1} \frac{s^m}{s^m + 1} ds.$$

To do so, we apply a change of variable $s = e^{-t/m}$. Then $ds = -\frac{1}{m}e^{-t/m}dt$ and, after re-arranging terms, the integral can be re-written as

$$I(m) = \int_0^\infty \frac{1 - e^{-t}}{t} \frac{1}{1 + e^t} \frac{t/m}{e^{t/m} - 1} dt. \quad (9.32)$$

It readily follows from [32, Theorem 3] that, for any $x > 0$ and $N \geq 1$, there exists a $\theta_{N,m} \in (0, 1)$, which might depend on m and N , such that

$$\frac{x}{e^x - 1} = 1 - \frac{x}{2} + \sum_{k=1}^{N-1} \frac{B_{2k}}{(2k)!} x^{2k} + \theta_{N,m} \frac{B_{2N}}{(2N)!} x^{2N}$$

Applying this observation to Equation (9.32) implies that

$$\begin{aligned} I(m) &= \int_0^\infty \frac{1 - e^{-t}}{t} \frac{1}{1 + e^t} \left(1 - \frac{t/m}{2} + \sum_{k=1}^{N-1} \frac{B_{2k}}{(2k)!} (t/m)^{2k} + \theta_{N,m} \frac{B_{2N}}{(2N)!} (t/m)^{2N} \right) dt \\ &= \ln\left(\frac{\pi}{2}\right) - \frac{2 \ln 2 - 1}{2m} + \sum_{k=1}^{N-1} \frac{B_{2k}}{(2k)!} \frac{1}{m^{2k}} I_2(k) + \theta_{N,m} \frac{B_{2N}}{(2N)!} \frac{1}{m^{2N}} I_2(N) \end{aligned}$$

where

$$I_2(k) = \int_0^\infty \frac{1 - e^{-t}}{1 + e^t} t^{2k-1} dt.$$

Using identity [33, 3.411.8], this integral can be evaluated as

$$\begin{aligned} I_2(k) &= \Gamma(2k) \left(\sum_{i=1}^{\infty} \frac{(-1)^{2k}}{i^{2k}} - \sum_{i=1}^{\infty} \frac{(-1)^{2k}}{(1+i)^{2k}} \right) \\ &= \Gamma(2k) \left(2 \sum_{i=1}^{\infty} \frac{(-1)^{2k}}{i^{2k}} - 1 \right) \\ &= \Gamma(2k) (2\eta(2k) - 1), \end{aligned}$$

where $\eta(s) = (1 - 2^{1-s})\zeta(s)$ is the Dirichlet eta function (with ζ the Riemann zeta function). We can thus write

$$I(m) = \ln\left(\frac{\pi}{2}\right) - \frac{2 \ln 2 - 1}{2m} + \sum_{k=1}^{N-1} \frac{B_{2k}}{2k} \frac{1}{m^{2k}} (2\eta(2k) - 1) + \theta_{N,m} \frac{B_{2N}}{2N} \frac{1}{m^{2N}} (2\eta(2N) - 1). \quad (9.33)$$

Now, by [34, Equation 9.11], there exists a $\sigma_{N,m} \in (0, 1)$, depending on m and N , such that

$$H_n = \ln m + \gamma + \frac{1}{2m} - \sum_{k=1}^{N-1} \frac{B_{2k}}{2k} \frac{1}{m^{2k}} - \sigma_{N,m} \frac{B_{2N}}{2N} \frac{1}{m^{2N}}, \quad (9.34)$$

where γ is the Euler-Mascheroni constant ($\gamma \approx 0.577$).

Equation (9.33) and Equation (9.34) imply that

$$H_n - I(m) = \ln\left(\frac{2m}{\pi}\right) + \gamma + \frac{\ln 2}{m} + \sum_{k=1}^{N-1} \frac{a_{2k}}{(2m/\pi)^{2k}} + C_{N,m} \frac{a_{2N}}{(2m/\pi)^{2N}} \quad (9.35)$$

with coefficients

$$a_{2k} := -\frac{B_{2k} 2\eta(2k)}{2k} \left(\frac{2}{\pi}\right)^{2k}$$

and constant

$$C_{N,m} = \frac{-\sigma_{N,m} - \theta_{N,m}(2\eta(2N) - 1)}{2\eta(2N)}.$$

The latter is derived from the equation

$$-\sigma_{N,m} \frac{B_{2N}}{2N} \frac{1}{m^{2N}} - \theta_{N,m} \frac{B_{2N}}{2N} \frac{1}{m^{2N}} (2\eta(2N) - 1) = C_{N,m} \frac{B_{2N}}{2N} \frac{1}{m^{2N}} 2\eta(2N)$$

Inserting Equation (9.35) into Equation (9.31) concludes the proof. \square

And in the special case where $N = 1$:

Corollary 9.4.8. *For any $m \in \mathbb{N}$, it holds that*

$$\sum_{k=1}^{m-1} \csc\left(\frac{\pi k}{m}\right) = \frac{2m}{\pi} \ln\left(\frac{2m}{\pi}\right) + \frac{2m}{\pi} \gamma + C_1 \frac{a_2}{2m} \pi. \quad (9.36)$$

This concludes our studies of sums of inverse sin's for now.

The following corollary is another useful building block on our mission to understand the expected complexity gain:

Corollary 9.4.9. *For any $m \in \mathbb{N}$, it is true that*

$$\sum_{i=1}^m \tan^{-1} |u(i)| + \sum_{i=0}^{m-1} \tan^{-1} |l(i)| \sim \frac{\pi}{4(2m+1)} \ln\left(\frac{\pi}{4(2m+1)}\right) + \frac{\pi}{4(2m+1)} \gamma + \mathcal{O}(m^{-2}). \quad (9.37)$$

Proof. It follows from the well known facts $\tan^{-1} x \sim x$ and $\sin x \sim x$ as $x \rightarrow 0$, that the sum

$$Z_m := \sum_{i=1}^m \tan^{-1} |\mathbf{u}(i)| + \sum_{i=0}^{m-1} \tan^{-1} |\mathbf{l}(i)|$$

behaves for $m \rightarrow \infty$ like

$$\begin{aligned} Z_m &= \sum_{i=1}^m \tan^{-1} \left(\frac{2 \sin^2 \left(\frac{\pi}{2} \frac{1}{2(2m+1)} \right)}{\sin \left(\frac{\pi}{2} \frac{2i}{2m+1} \right)} \right) + \sum_{i=0}^{m-1} \tan^{-1} \left(\frac{2 \sin^2 \left(\frac{\pi}{2} \frac{1}{2(2m+1)} \right)}{\sin \left(\frac{\pi}{2} \frac{2i+1}{2m+1} \right)} \right) \\ &= \sum_{i=1}^{2m} \tan^{-1} \left(\frac{2 \sin^2 \left(\frac{\pi}{2} \frac{1}{2(2m+1)} \right)}{\sin \left(\frac{\pi}{2} \frac{i}{2m+1} \right)} \right) \\ &\sim 2 \left(\frac{\pi}{2} \frac{1}{2(2m+1)} \right)^2 \sum_{i=1}^{2m} \frac{1}{\sin \left(\frac{\pi}{2} \frac{i}{2m+1} \right)} \\ &\stackrel{9.4.4}{=} \left(\frac{\pi}{2} \frac{1}{2(2m+1)} \right)^2 \left(\sum_{i=1}^{4m+1} \frac{1}{\sin \left(\frac{\pi}{4m+2} \frac{i}{2} \right)} - 1 \right). \end{aligned}$$

Using Corollary 9.4.8, the sum can be approximated as

$$\sum_{i=1}^{4m+1} \frac{1}{\sin \left(\frac{\pi i}{4m+2} \right)} = \frac{2(4m+2)}{\pi} \ln \left(\frac{2(4m+2)}{\pi} \right) + \frac{2(4m+2)}{\pi} \gamma + \mathcal{O}(m^{-1})$$

and thus

$$Z_m \sim \frac{\pi}{4(2m+1)} \ln \left(\frac{4(2m+1)}{\pi} \right) + \frac{\pi}{4(2m+1)} \gamma + \mathcal{O}(m^{-2}).$$

□

Remark 9.4.10. Note that one could drop the $\mathcal{O}(m^{-2})$ in Equation (9.37) since

$$\frac{\pi}{4(2m+1)} \ln \left(\frac{\pi}{4(2m+1)} \right) + \frac{\pi}{4(2m+1)} \gamma + \mathcal{O}(m^{-2}) \sim \frac{\pi}{4(2m+1)} \ln \left(\frac{\pi}{4(2m+1)} \right) + \frac{\pi}{4(2m+1)} \gamma.$$

However, we keep the non-leading terms for higher resolution (even though we dropped resolution when approximating $\sin x \sim x$ and $\tan^{-1} x \sim x$ in the proof of Corollary 9.4.9).

Having developed the above machinery for sums of cotangents, we now turn our attention to computing the expected number $\mathbb{E}[S]$ of points in the upper convex hull of $P_1 \diamond N_1$.

The following lemma provides a closed-form expression for the conditional expectations:

Proposition 9.4.11. *The following identities hold for the conditional expectations of S :*

$$\mathbb{E}[S \mid w > 0, b > 0] = 2n + 2 + \frac{2}{\pi} \sum_{i=1}^n \tan^{-1} |\mathbf{u}(i)| \quad (9.38)$$

$$\mathbb{E}[S \mid w > 0, b < 0] = n + 2 + \frac{4}{\pi} \sum_{i=0}^{n-1} \tan^{-1} |\mathbf{l}(i)| \quad (9.39)$$

$$\mathbb{E}[S \mid w < 0, b > 0] = 2n + 2 + \frac{2}{\pi} \sum_{i=0}^{n-1} \tan^{-1} |\mathbf{l}(i)| \quad (9.40)$$

$$\mathbb{E}[S \mid w < 0, b < 0] = n + 2 + \frac{4}{\pi} \sum_{i=1}^n \tan^{-1} |\mathbf{u}(i)|. \quad (9.41)$$

Proof. We start with Equation (9.38). Writing $S' := S - (2n + 2)$, it follows from Theorem 9.3.6 that

$$\begin{aligned} \mathbb{E}[S' \mid w > 0, b > 0] &= \frac{2}{\pi} \left(\sum_{i=1}^{n-1} i (\tan^{-1} |\mathbf{u}(i)| - \tan^{-1} |\mathbf{u}(i+1)|) + n \tan^{-1} |\mathbf{u}(n)| \right) \\ &= \frac{2}{\pi} \left(\sum_{i=1}^n i \tan^{-1} |\mathbf{u}(i)| - \sum_{i=1}^{n-1} i \tan^{-1} |\mathbf{u}(i+1)| \right) \\ &= \frac{2}{\pi} \left(\sum_{i=1}^n i \tan^{-1} |\mathbf{u}(i)| - \sum_{i=2}^n (i-1) \tan^{-1} |\mathbf{u}(i)| \right) \\ &= \frac{2}{\pi} \left(\sum_{i=2}^n \tan^{-1} |\mathbf{u}(i)| + \tan^{-1} |\mathbf{u}(1)| \right) \\ &= \frac{2}{\pi} \sum_{i=1}^n \tan^{-1} |\mathbf{u}(i)|. \end{aligned}$$

This shows Equation (9.38). Analogously, one can prove Equation (9.40), this time using Theorem 9.3.11:

$$\begin{aligned} \mathbb{E}[S' \mid w < 0, b > 0] &= \frac{2}{\pi} \left(\sum_{i=1}^{n-1} i (\tan^{-1} |\mathbf{l}(i-1)| - \tan^{-1} |\mathbf{l}(i)|) + n \tan^{-1} |\mathbf{l}(n-1)| \right) \\ &= \frac{2}{\pi} \left(\sum_{i=1}^n i \tan^{-1} |\mathbf{l}(i-1)| - \sum_{i=1}^{n-1} i \tan^{-1} |\mathbf{l}(i)| \right) \\ &= \frac{2}{\pi} \left(\sum_{i=0}^{n-1} (i+1) \tan^{-1} |\mathbf{l}(i)| - \sum_{i=1}^{n-1} i \tan^{-1} |\mathbf{l}(i)| \right) \\ &= \frac{2}{\pi} \left(\sum_{i=1}^{n-1} \tan^{-1} |\mathbf{l}(i)| + \tan^{-1} |\mathbf{l}(0)| \right) \\ &= \frac{2}{\pi} \sum_{i=0}^{n-1} \tan^{-1} |\mathbf{l}(i)|. \end{aligned}$$

Equations (9.39) and (9.41) follow analogously to Equation (9.38) and (9.40) from Theorems 9.3.7 and 9.3.12. \square

Finally, the law of total expectation allows stitching together the conditional expectations to obtain the unconditional expectation of S :

Theorem 9.4.12. *The expected number of points in the upper convex hull of $P_1 \diamond N_1$ is given by*

$$\mathbb{E}[S] \sim \frac{3}{2}n + 2 + \frac{3}{8(2n+1)} \ln \left(\frac{4(2n+1)}{\pi} \right) + \frac{3}{8(2n+1)}\gamma + \mathcal{O}(n^{-2}). \quad (9.42)$$

Proof. Following the law of total expectation, $\mathbb{E}[S]$ splits into a sum of conditional expectations,

$$\mathbb{E}[S] = \frac{1}{4}(\zeta_1 + \zeta_2),$$

where

$$\begin{aligned} \zeta_1 &:= \mathbb{E}[S \mid w > 0, b > 0] + \mathbb{E}[S \mid w < 0, b > 0] \\ \zeta_2 &:= \mathbb{E}[S \mid w > 0, b < 0] + \mathbb{E}[S \mid w < 0, b < 0]. \end{aligned}$$

As a next step, we study ζ_1 and ζ_2 more closely. By Proposition 9.4.11, we can reformulate

$$\zeta_1 = 4n + 4 + \frac{2}{\pi} \left(\sum_{i=1}^n \tan^{-1} |\mathbf{u}(i)| + \sum_{i=0}^{n-1} \tan^{-1} |\mathbf{l}(i)| \right).$$

Similarly, it holds that

$$\zeta_2 = 2n + 4 + \frac{4}{\pi} \left(\sum_{i=1}^n \tan^{-1} |\mathbf{u}(i)| + \sum_{i=0}^{n-1} \tan^{-1} |\mathbf{l}(i)| \right),$$

so that

$$\zeta_1 + \zeta_2 = 6n + 8 + \frac{6}{\pi} \left(\sum_{i=1}^n \tan^{-1} |\mathbf{u}(i)| + \sum_{i=0}^{n-1} \tan^{-1} |\mathbf{l}(i-1)| \right).$$

By Corollary 9.4.9, this implies that

$$\zeta_1 + \zeta_2 \sim 6n + 8 + \frac{6}{\pi} \left(\frac{\pi}{4(2n+1)} \ln \left(\frac{4(2n+1)}{\pi} \right) + \frac{\pi}{4(2n+1)}\gamma + \mathcal{O}(n^{-2}) \right).$$

This concludes the proof. \square

A natural question arising from Theorem 9.4.12 is the following: how does the expected marginal complexity

$$\mathcal{C}^\uparrow := \mathbb{E}[S] - (2n + 1),$$

which captures the gain in d -cells through the additional layer, behave for large n ? The answer is now a simple corollary:

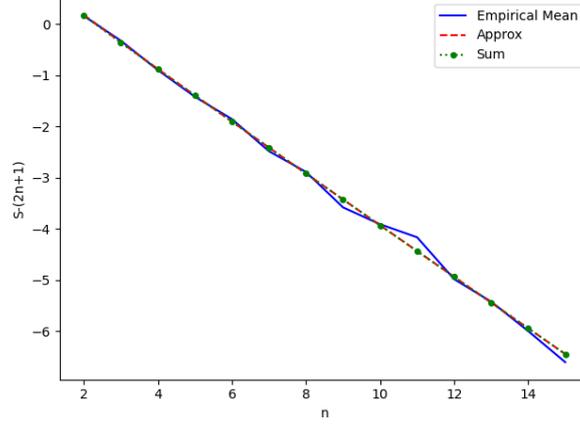


Figure 9.11: The expected marginal complexity C^\uparrow as a function of n . Blue line is the empirical mean, averaged over 1000 random initializations. Dashed line corresponds to the leading terms approximation given in Corollary 9.4.13. Green dotted line corresponds to the exact sum derived from Theorem 9.3.14.

Corollary 9.4.13. *The expected marginal complexity behaves like*

$$C^\uparrow \sim -\frac{1}{2}n + 1 + \frac{3}{8(2n+1)} \ln\left(\frac{4(2n+1)}{\pi}\right) + \frac{3}{8(2n+1)}\gamma + \mathcal{O}(n^{-2}) \quad (9.43)$$

In words, we can expect the number of d -cells to decrease linearly in n . This result is confirmed in Figure 9.11, which plots the empirical mean, the exact mean using Theorem 9.3.14 as well as the approximation in Corollary 9.4.13.

Corollary 9.4.13 says that the number of d -cells is expected to decrease under the additional ReLU layer. However, by Corollary 7.1.8, the d -cells can be finer than affine regions. The following corollary establishes a similar result for the number of affine regions:

Corollary 9.4.14. *The number of affine regions is expected to decrease under the additional random ReLU layer.*

Proof. By Corollary 7.1.8, $S = |\mathcal{U}^*(P_1 \diamond N_1)|$ is an upper bound for the number of affine regions defined by $\mathcal{Q}(P_1) - \mathcal{Q}(N_1)$ (see also Remark 9.3.4). By Proposition 9.4.1, the number of affine regions, like the number of d -cells, defined by $\mathcal{Q}(P_0) - \mathcal{Q}(N_0)$ is also given by $2n + 1$. \square

What is the take-away from this result? In the beginning of this chapter, we conjectured, inspired by the dual representation of Telgarsky's network derived in Chapter 8, that the network $(P_0, N_0)^3$ lives in the exponential complexity regime.

In particular, since $\mathcal{U}(P_0 \diamond N_0) = \mathcal{O}(n)$ by Proposition 9.2.1, this means that $P_0 \diamond N_0$ can be written using $\log n$ narrow ReLU layers.

³or a close approximation thereof, as it remains to show that N_0 and P_0 can be placed on the sphere by piecewise linear activation functions

Corollary 9.4.14 says that adding one random ReLU layer after these $\log n$ deterministic layers is expected to *decrease* the number of affine regions. This confirms our claim that the exponential complexity regime is unstable/sharp. The random weights transition the network to the second regime.

Chapter 10

Heuristic Arguments in Dual Space

In this chapter, we study the volume hypothesis using two additional toy examples. Compared to the problem of summing circles discussed in the previous chapter, the problems here focus solely on understanding the complexity of the upper convex hull. To do so, these problems are abstract and take place solely in dual space. In particular, the randomness over dual points does not arise from weights and biases but is instead assumed to follow a known distribution; uniform in Section 10.1 and Gaussian in Section 10.2.

In Section 10.1, we show that the probability of increasing the size of an upper convex hull of size n by adding one point uniformly at random diminishes for large n . In Section 10.2, we establish a similar result for two sets of i.i.d. Gaussian red and blue random points and edges between them.

10.1 Large Upper Convex Hulls

In this section, we show (in a specific setup) that the upper convex hull, upon adding a new point at random, grows more slowly if it is already largely occupied.

This is a simplification, since the real dynamics of a neural network are such that the whole representation changes from layer to layer. Nevertheless, this approach offers an interesting perspective on the dual representation and provides another heuristic to think about the volume hypothesis: large convex hulls, corresponding to complex networks, do not have a lot of space left to grow.

In our specific case, we define

$$Q := \{(x, y) \in \mathbb{R}^2 \mid 0 \leq x, y \leq 1\},$$

to be the two-dimensional unit cube with lower left corner at the origin and $\mathcal{X}_n \subseteq S$ to be a set of $n + 1$ equidistant points on the part of the sphere contained in the first quadrant,

$$\mathcal{X}_n = \left\{ \left(\cos \left(\frac{\pi i}{2n} \right), \sin \left(\frac{\pi i}{2n} \right) \right) \mid i \in \mathbb{N}_0, 0 \leq i \leq n \right\}.$$

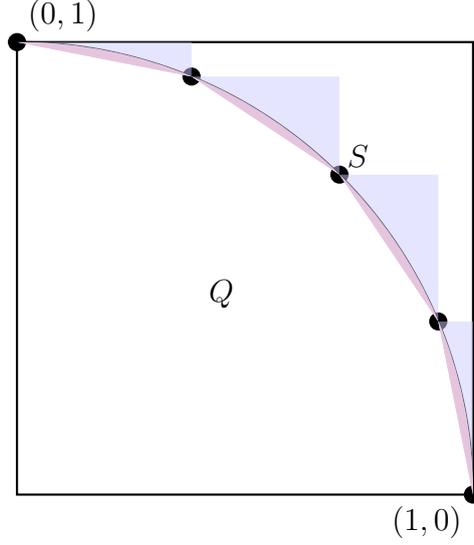


Figure 10.1: Example drawing for $n = 4$. The blue shaded areas are the ω_i^n , adding the red areas gives the Δ_i^n .

Note that $\mathcal{U}^*(\mathcal{X}_n) = \mathcal{X}_n$, i.e., that \mathcal{X}_n already contains all the vertices of its upper convex hull. We are interested in the following question:

Upon sampling a point $y \sim U(Q)$ uniformly at random from the cube, what is the probability that y will increase the size of the upper convex hull? That is, what is the probability that $|\mathcal{U}^*(\mathcal{X}_n \cup \{y\})| > |\mathcal{U}^*(\mathcal{X}_n)|$?

To answer this question, note that the size of the upper convex hull can only grow if y falls into the set Ω defined as

$$\Omega := \cup_{i=0}^{n-1} \omega_i^n$$

where

$$\omega_i^n = \Delta_i^n \cap \{x \in \mathbb{R}^2 \mid \|x\|_2^2 > 1\}$$

with Δ_i^n the simplex on points $(\cos(\frac{\pi i}{2n}), \sin(\frac{\pi i}{2n}))$, $(\cos(\frac{\pi(i+1)}{2n}), \sin(\frac{\pi(i+1)}{2n}))$ and $(\cos(\frac{\pi i}{2n}), \sin(\frac{\pi(i+1)}{2n}))$ (see Figure 10.1).

Since y is picked uniformly at random from Q , it follows that the probability of y increasing the size of the upper convex hull is given by

$$\mathbb{P}(|\mathcal{U}^*(\mathcal{X}_n \cup \{y\})| > |\mathcal{U}^*(\mathcal{X}_n)|) = \frac{V(\Omega)}{V(Q)} = V(\Omega), \quad (10.1)$$

where $V(\Omega)$ and $V(Q)$ are the volume of Ω and Q , respectively.

In the following, we will compute an upper bound for this probability, using the fact that

$$\Omega \subseteq \cup_{i=1}^{n-1} \Delta_i^n =: \Delta^n. \quad (10.2)$$

Proposition 10.1.1. *The volume of Δ^n is given by*

$$V(\Delta^n) = \frac{1 - \cos(\frac{\pi}{2n})}{2 \sin(\frac{\pi}{2n})}.$$

Proof. The volume of the i 'th triangle is given by

$$\begin{aligned}
V(\Delta_i^n) &= \frac{1}{2} \left(\cos \left(\frac{\pi i}{2n} \right) - \cos \left(\frac{\pi(i+1)}{2n} \right) \right) \cdot \left(\sin \left(\frac{\pi(i+1)}{2n} \right) - \sin \left(\frac{\pi i}{2n} \right) \right) \\
&= \frac{1}{2} \left(\cos \left(\frac{\pi i}{2n} \right) \sin \left(\frac{\pi(i+1)}{2n} \right) - \cos \left(\frac{\pi(i+1)}{2n} \right) \sin \left(\frac{\pi i}{2n} \right) \right) \\
&\quad - \cos \left(\frac{\pi(i+1)}{2n} \right) \sin \left(\frac{\pi(i+1)}{2n} \right) + \cos \left(\frac{\pi i}{2n} \right) \sin \left(\frac{\pi i}{2n} \right) \\
&\stackrel{\text{A.0.5}}{=} \frac{1}{4} \left[\sin \left(\frac{\pi(2i+1)}{2n} \right) + \sin \left(\frac{\pi}{2n} \right) - \sin \left(\frac{\pi(2i)}{2n} \right) - \sin(0) - \right. \\
&\quad \left. - \sin \left(\frac{\pi(2(i+1))}{2n} \right) - \sin(0) + \sin \left(\frac{\pi(2i+1)}{2n} \right) - \sin \left(\frac{\pi}{2n} \right) \right] \\
&= \frac{1}{4} \left[2 \sin \left(\frac{\pi(2i+1)}{2n} \right) - \sin \left(\frac{\pi(2i)}{2n} \right) - \sin \left(\frac{\pi(2(i+1))}{2n} \right) \right] \\
&= \frac{1}{4} \left[2 \sin \left(\frac{\pi(2i+1)}{2n} \right) - \sin \left(\frac{\pi i}{n} \right) - \sin \left(\frac{\pi(i+1)}{n} \right) \right].
\end{aligned}$$

Consequently, the total area of all triangles is

$$V(\Delta^n) = \sum_{i=0}^{n-1} V(\Delta_i^n) = \frac{1}{4} \left[2 \sum_{i=0}^{n-1} \sin \left(\frac{\pi(2i+1)}{2n} \right) - \sum_{i=0}^{n-1} \sin \left(\frac{\pi i}{n} \right) - \sum_{i=0}^{n-1} \sin \left(\frac{\pi(i+1)}{n} \right) \right].$$

We compute each sum separately using Lemma A.0.6:

$$\begin{aligned}
\sum_{i=0}^{n-1} \sin \left(\frac{\pi(2i+1)}{2n} \right) &= \sum_{i=0}^{n-1} \sin \left(\frac{\pi}{2n} + \frac{\pi i}{n} \right) \\
&= \frac{\sin \left(n \cdot \frac{\pi}{2n} \right)}{\sin \left(\frac{\pi}{2n} \right)} \sin \left(\frac{\pi}{2n} + (n-1) \frac{\pi}{2n} \right) \\
&= \frac{1}{\sin \left(\frac{\pi}{2n} \right)}.
\end{aligned}$$

Next,

$$\begin{aligned}
\sum_{i=0}^{n-1} \sin \left(\frac{\pi i}{n} \right) &= \sum_{i=0}^{n-1} \sin \left(\frac{\pi i}{n} \right) \\
&= \frac{\sin \left(n \cdot \frac{\pi}{2n} \right)}{\sin \left(\frac{\pi}{2n} \right)} \sin \left((n-1) \frac{\pi}{2n} \right) \\
&= \frac{\sin \left(\frac{\pi}{2} - \frac{\pi}{2n} \right)}{\sin \left(\frac{\pi}{2n} \right)} \\
&= \frac{\cos \left(\frac{\pi}{2n} \right)}{\sin \left(\frac{\pi}{2n} \right)} \\
&= \cot \left(\frac{\pi}{2n} \right).
\end{aligned}$$

Finally,

$$\begin{aligned}
\sum_{i=0}^{n-1} \sin\left(\pi \frac{i+1}{n}\right) &= \sum_{i=0}^{n-1} \sin\left(\frac{\pi}{n} + \frac{\pi i}{n}\right) \\
&= \frac{\sin\left(n \cdot \frac{\pi}{2n}\right)}{\sin\left(\frac{\pi}{2n}\right)} \sin\left(\frac{\pi}{n} + (n-1)\frac{\pi}{2n}\right) \\
&= \frac{\sin\left(\frac{\pi}{2} + \frac{\pi}{2n}\right)}{\sin\left(\frac{\pi}{2n}\right)} \\
&= \frac{\cos\left(\frac{\pi}{2n}\right)}{\sin\left(\frac{\pi}{2n}\right)} \\
&= \cot\left(\frac{\pi}{2n}\right).
\end{aligned}$$

We conclude that

$$V(\Delta^n) = \frac{1}{4} \left(2 \frac{1}{\sin\left(\frac{\pi}{2n}\right)} - 2 \cot\left(\frac{\pi}{2n}\right) \right).$$

The claim then follows from a simple reformulation. \square

The following theorem uses this formula for the volume $V(\Delta^n)$ to upper bound the probability of enlarging the upper convex hull, and shows that it vanishes for large n .

Theorem 10.1.2. *The probability of y increasing the size of the upper convex hull is upper bounded by*

$$\mathbb{P}(|\mathcal{U}^*(\mathcal{X}_n \cup \{y\})| > |\mathcal{U}^*(\mathcal{X}_n)|) \leq \frac{1 - \cos\left(\frac{\pi}{2n}\right)}{\sin\left(\frac{\pi}{2n}\right)}.$$

Furthermore, this probability tends to zero for large n :

$$\lim_{n \rightarrow \infty} \mathbb{P}(|\mathcal{U}^*(\mathcal{X}_n \cup \{y\})| > |\mathcal{U}^*(\mathcal{X}_n)|) = 0.$$

Proof. By construction and Proposition 10.1.1,

$$\mathbb{P}(|\mathcal{U}^*(\mathcal{X}_n \cup \{y\})| > |\mathcal{U}^*(\mathcal{X}_n)|) \leq V(\Delta^n) = \frac{1 - \cos\left(\frac{\pi}{2n}\right)}{\sin\left(\frac{\pi}{2n}\right)}.$$

This shows the first claim.

For the second claim, apply L'Hospital's rule to conclude that

$$\begin{aligned}
\lim_{n \rightarrow \infty} V(\Delta^n) &= \lim_{n \rightarrow \infty} \frac{1 - \cos\left(\frac{\pi}{2n}\right)}{\sin\left(\frac{\pi}{2n}\right)} \\
&= \lim_{n \rightarrow \infty} \frac{\frac{\pi}{2n^2} \sin\left(\frac{\pi}{2n}\right)}{-\frac{\pi}{2n^2} \cos\left(\frac{\pi}{2n}\right)} \\
&= - \lim_{n \rightarrow \infty} \frac{\sin\left(\frac{\pi}{2n}\right)}{\cos\left(\frac{\pi}{2n}\right)} \\
&= 0.
\end{aligned}$$

\square

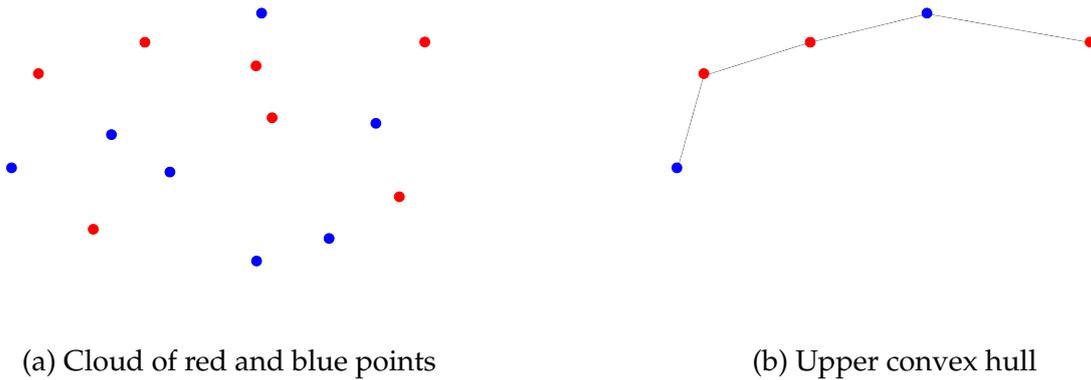


Figure 10.2: An example of red and blue points in \mathbb{R}^2 with $\alpha = 7$ and $\beta = 7$. Their upper convex hull can be thought of as a chain of red and blue nobs.

This theorem shows that the probability of enlarging the upper convex hull vanishes as more points are already in the upper convex hull. As described in the introduction to this section, this result provides another heuristic to think about the volume hypothesis. In the setting of neural networks, it reads that there are less ways to grow the complexity of already complex networks than that of simple ones.

10.2 Gaussian Dual Points

In this section, we provide further abstract considerations to better understand large upper convex hulls, this time in relation to decision boundary complexity in the sense of Proposition 6.2.2. In particular, given two sets $\mathcal{R}, \mathcal{B} \subseteq \mathbb{R}^2$, we are interested in edges inside $\mathcal{U}(\mathcal{R} \cup \mathcal{B})$, which contain points from both \mathcal{R} and \mathcal{B} .

We start with some basic definitions. Let $\mathcal{R}, \mathcal{B} \subseteq \mathbb{R}^2$ be two non-empty sets containing i.i.d. multivariate uncorrelated Gaussian random variables, with size $|\mathcal{R}| = \alpha$, $|\mathcal{B}| = \beta$ and $n := \alpha + \beta$. The random variables belonging to \mathcal{R} are called *red points*, while those belonging to \mathcal{B} are called *blue points*.

The set containing all points is called $\mathcal{X} := \mathcal{R} \cup \mathcal{B}$. Enumerate its entries $\mathcal{X} = \{P_1, \dots, P_n\}$, where each $P_i \stackrel{\text{iid}}{\sim} \mathcal{N}(0, 1)$.

We call an edge in $\mathcal{C}(\mathcal{X})$ *good* if it contains both a red and a blue point. In accordance with Proposition 6.2.2, we are interested in counting the good edges in $\mathcal{U}(\mathcal{X})$.

To do so, one first needs to understand the probability that an edge (P_i, P_j) is contained in just the convex hull of \mathcal{X} . Define the corresponding event

$$E_{ij} := \text{edge } (P_i, P_j) \text{ is contained in } \mathcal{C}(\mathcal{X}).$$

The following proposition provides the probability for this event.

Proposition 10.2.1 (Boundary Edges [35, Satz 4]). *The probability that an edge (P_i, P_j) is*

is contained in the convex hull $\mathcal{C}(\mathcal{X})$ is given by

$$\mathbb{P}(E_{ij} = 1) = \frac{2}{\sqrt{\pi}} \int_0^\infty (\Phi(p)^{n-2} + (1 - \Phi(p))^{n-2}) e^{-p^2} dp \quad (10.3)$$

where

$$\Phi(p) = \frac{1}{\sqrt{2\pi}} \int_{-\infty}^p e^{-\frac{u^2}{2}} du.$$

For large n , this probability satisfies

$$\mathbb{P}(E_{ij} = 1) \sim \frac{4\sqrt{2\pi \log n}}{n^2}.$$

Proof. We start by showing the non-asymptotic identity. Since the points in \mathcal{X} are i.i.d., it suffices to compute $\mathbb{P}(E_{12} = 1)$.

Let L be the line going through P_1 and P_2 and write S_1 and S_2 for the two closed half-spaces induced by L . Then $E_{ij} = 1$ if and only if all the other points lie either in S_1 or S_2 . Defining

$$A := \{P_k \in S_1 \forall k \geq 3 \text{ or } P_k \in S_2 \forall k \geq 3\}$$

allows writing

$$\mathbb{P}(E_{12} = 1) = \mathbb{P}(A) \quad (10.4)$$

$$= \int \mathbb{P}(A \mid P_1 = p_1, P_2 = p_2) p_{P_2}(p_2) p_{P_1}(p_1) dp_1 dp_2, \quad (10.5)$$

where p_{P_1} and p_{P_2} are the densities of P_1 and P_2 (both Gaussian by assumption).

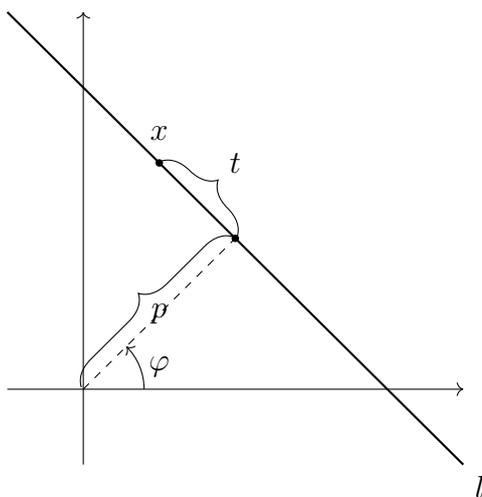


Figure 10.3: Explanation of line coordinates for a line l (inspired by [36, Figure 4.1]).

Since the P_i are assumed to be i.i.d.,

$$\begin{aligned}\mathbb{P}(A) &= \mathbb{P}(P_k \in S_1 \forall k \geq 3) + \mathbb{P}(P_k \in S_2 \forall k \geq 3) \\ &= \mathbb{P}(P_3 \in S_1)^{n-2} + \mathbb{P}(P_3 \notin S_1)^{n-2} \\ &= \mathbb{P}(P_3 \in S_1)^{n-2} + (1 - \mathbb{P}(P_3 \in S_1))^{n-2}.\end{aligned}$$

Assume now that $P_1 = p_1 \in \mathbb{R}^2$, $P_2 = p_2 \in \mathbb{R}^2$ are realized, and write l for the line connecting p_1 and p_2 . Similarly, write s_1, s_2 for half-spaces induced by l . Analogously to the last computation, one can show that

$$\mathbb{P}(A \mid P_1 = p_1, P_2 = p_2) = \mathbb{P}(P_3 \in s_1)^{n-2} + (1 - \mathbb{P}(P_3 \in s_1))^{n-2}. \quad (10.6)$$

In order to study the integral resulting from inserting Equation (10.6) into the integral in Equation (10.5), we introduce line coordinates [36] to express points on l . In these coordinates, any point $x \in l$ on the line can be written as

$$x = p \begin{pmatrix} \cos \varphi \\ \sin \varphi \end{pmatrix} + t \begin{pmatrix} -\sin \varphi \\ \cos \varphi \end{pmatrix}$$

where p is the the shortest distance from l to the origin, t is the distance from x to the intersection of l with the line determined by the shortest distance to the origin, and φ is the angle between the positive first axis and r (see Figure 10.3). A simple calculation shows that

$$\|x\|^2 = p^2 + t^2. \quad (10.7)$$

Furthermore, by [36, Equation (4.2)], the volume form translates like

$$dp_1 dp_2 = dp d\varphi dt_1 dt_2 |t_1 - t_2|, \quad (10.8)$$

where (p, φ, t_i) are the line-coordinates for p_i on l for both $i = 1, 2$.

We come back to computing $\mathbb{P}(P_3 \in s_1)$, now using line coordinates. Without loss of generality, assume that the half-space s_1 realized by l does not contain the origin. Then, by the isotropy of the Gaussian distribution, we may rotate the line l around the origin until it is parallel to the first axis. In this case,

$$\mathbb{P}(P_3 \in s_1) = \mathbb{P}((P_3)_2 \geq p) \quad (10.9)$$

$$= \frac{1}{2\pi} \int_{\mathbb{R}} \int_p^{+\infty} e^{-\frac{x^2+y^2}{2}} dx dy \quad (10.10)$$

$$= \frac{1}{\sqrt{2\pi}} \int_p^{+\infty} e^{-\frac{y^2}{2}} dy \quad (10.11)$$

$$= 1 - \Phi(p). \quad (10.12)$$

Finally, using Equation (10.7), we deduce that

$$\frac{1}{2\pi} e^{-\frac{\|p_i\|^2}{2}} = \frac{1}{2\pi} e^{-\frac{p^2+t_i^2}{2}} \quad (10.13)$$

for both $i = 1, 2$.

Inserting Equations (10.6), (10.8), (10.12) and (10.13) into Equation (10.5) gives

$$\begin{aligned}\mathbb{P}(E_{ij}) &= \frac{1}{(2\pi)^2} \int_0^{2\pi} d\varphi \iiint (\Phi(p)^{n-2} + (1 - \Phi(p))^{n-2}) |t_1 - t_1| e^{-p^2 - \frac{t_1^2 + t_2^2}{2}} dt_1 dt_2 dp \\ &= \frac{1}{2\pi} \iiint (\Phi(p)^{n-2} + (1 - \Phi(p))^{n-2}) |t_1 - t_1| e^{-p^2 - \frac{t_1^2 + t_2^2}{2}} dt_1 dt_2 dp.\end{aligned}$$

Next we claim that

$$\iint |t_1 - t_2| e^{-\frac{t_1^2 + t_2^2}{2}} dt_1 dt_2 = 4\sqrt{\pi}, \quad (10.14)$$

which would show Equation (10.3) and conclude the proof of the first claim of the Proposition. Indeed, define the transformation

$$(u, v) = \left(\frac{t_1 - t_2}{\sqrt{2}}, \frac{t_1 + t_2}{\sqrt{2}} \right).$$

Simple calculations reveal that

$$\begin{aligned}(t_1, t_2) &= \left(\frac{u + v}{\sqrt{2}}, \frac{v - u}{\sqrt{2}} \right) \\ t_1^2 + t_2^2 &= u^2 + v^2 \\ dt_1 dt_2 &= du dv \\ t_1 - t_2 &= \sqrt{2}u\end{aligned}$$

and therefore

$$\begin{aligned}\iint |t_1 - t_2| e^{-\frac{t_1^2 + t_2^2}{2}} dt_1 dt_2 &= \sqrt{2} \int du |u| e^{-\frac{u^2}{2}} \int dv e^{-\frac{v^2}{2}} \\ &= \sqrt{2} \sqrt{2\pi} 2 \\ &= 4\sqrt{\pi}.\end{aligned}$$

This shows the first claim.

For the asymptotic behavior, see the second part of the proof of [35, Satz 4]. \square

With the probability that any specific edge is in the boundary of the convex hull at hand, the more specific case concerning the upper convex hull follows readily. Define the corresponding event

$$U_{ij} := \text{edge}(P_i, P_j) \text{ is contained in } \mathcal{U}(\mathcal{X}).$$

Then the total number of good edges is given by

$$U := \sum_{\substack{P_i \in \mathcal{R}, \\ P_j \in \mathcal{B}}} U_{ij}.$$

The symmetry of the problem implies that the probability of any specific edge being in the upper convex hull is

$$\mathbb{P}(U_{12} = 1) = \mathbb{P}(U_{ij} = 1) = \frac{1}{2}\mathbb{P}(E_{ij} = 1).$$

Putting these results together, the following theorem provides the expected number of good edges:

Theorem 10.2.2. *The expected number of edges in $\mathcal{X} = \mathcal{R} \cup \mathcal{B}$ containing points from both \mathcal{R} and \mathcal{B} satisfies*

$$\mathbb{E}[U] \sim \frac{\alpha(n - \alpha)}{2} \frac{4\sqrt{2\pi \log n}}{n^2}, \quad n \rightarrow \infty. \quad (10.15)$$

Proof. Clearly

$$\mathbb{E}[U] = \sum_{\substack{P_i \in \mathcal{R}, \\ P_j \in \mathcal{B}}} \mathbb{E}[U_{ij}] = \frac{1}{2} \sum_{\substack{P_i \in \mathcal{R}, \\ P_j \in \mathcal{B}}} \mathbb{P}(E_{ij} = 1).$$

The claim then follows from Proposition 10.2.1. □

For the rest of section, we discuss the implication of Theorem 10.2.2.

Given a fixed number of points, the expected number of good edges is maximized for $\alpha = n/2$, in which case

$$\mathbb{E}[U] \sim \frac{\sqrt{2\pi \log n}}{2}.$$

In particular, the marginal gain diminishes for large n like

$$\frac{d}{dn} \sqrt{\log n} = \frac{1}{2n \log n}.$$

While this result is not as strong as the one in Corollary 9.4.13 (which derives negative marginal complexity, although in a completely different setting), this shows that the marginal gain obtained by including one more point goes to zero.

This concludes our discussion of the first extreme case. In the second extreme case, Equation (10.15) is minimal for $\alpha = 1$ or $\alpha = n - 1$, and

$$\mathbb{E}[U] \sim \frac{2\sqrt{2\pi \log n}}{n}.$$

In particular, the marginal gain obtained by adding one point is negative in this case.

Interpolating between these two cases shows that, for any ratio of red and blue points, the marginal gain in the expected number of good edges transitions from positive but diminishing to negative.

Chapter 11

Conclusion

The first part of this thesis summarizes and elaborates on existing work on the dual representation of fully connected feedforward ReLU networks. It provides two equivalent perspectives to think about this representation, one based on tropical affine and the other on tropical geometry.

Using the dual representation, we fill gaps in existing arguments to rigorously derive duality results which translate complexity measures to dual space. In particular, the number of affine regions defined by a network $\mathcal{N} = \mathcal{Q}(P) - \mathcal{Q}(N): \mathbb{R}^d \rightarrow \mathbb{R}$ corresponds to the number of vertices in the upper convex hull of $P \diamond N$, and the number of boundary pieces of a binary classification network equals the number of edges in $\mathcal{U}(P \cup N)$ containing points from both P and N .

In the second part of this work, we use the previously derived duality results to provide evidence for the volume hypothesis. This hypothesis was formulated by Chiang et al. [6] to explain why overparameterized models trained using a Guess & Check algorithm generalize well, despite not being implicitly regularized by first order optimization techniques. Chiang et al. attribute this observation solely to the topology of the loss landscape, arguing that well generalizing minima have a larger volume in parameter space (i.e., are more flat). Throughout this work, we use simplicity as a proxy for generalization capabilities. Crucially, we present a novel way to think about generalization capabilities of deep ReLU network using affine geometry.

In different low-dimensional settings, we provide evidence for the volume hypothesis by showing the existence of a simplicity bias: exponentially complex networks are unstable, as small changes to the network quickly lead to sub-exponential complexity.

To better understand the dual representation, we derive it for Telgarsky's sawtooth function in Chapter 8. Telgarsky achieves a number of affine regions which is exponential in the depth of the network, close to the theoretical maximum. We show, by randomizing the last layers of the sawtooth network, that his maximum is unstable, concluding that the corresponding minimum in the loss-landscape is sharp. In particular, exponential complexity seems to require a careful choice of weights, which is again evidence for the volume hypothesis.

In Chapter 9, we replace numerical simulations with mathematical analysis, studying how the complexity of a deterministic network, inspired by Telgarsky’s network and conjectured to have exponential complexity, changes when appending a ReLU layer with Gaussian weight and bias. This approach shows that the expected marginal gain, which is the difference in the number of affine regions between the larger and smaller networks, decreases as the size of the deterministic network increases. This observation further supports the volume hypothesis, demonstrating that the number of affine regions is less stable the more complex the network is.

Finally, these results are illustrated in two additional toy settings. Formulated purely in dual space, these problems are simplified heuristics, and assume the dual points follow a known distribution, such as a Gaussian.

Future Work

We see two main tracks for future work. The first should deepen the understanding of the dual representation of ReLU networks. The second should further explore the volume hypothesis.

On the first track, Conjecture 7.1.10 states that, for a random network $\mathcal{Q}(P) - \mathcal{Q}(N)$, the d -cells are almost surely the affine regions. A proof of this conjecture would allow working with affine regions, not just activation regions or other refinements like is often done in existing literature [14]. Further studies of the graph-structure inducing $\mathcal{U}^*(P \diamond N) / \sim$ could allow approaching the problem from the perspective of graph theory.

We would also be interested in better understanding how sensitive the dual representation is to small changes in the networks weights and biases. Potential results could involve bounds on the displacement of points in P and N given a small change in the weights in biases, or they could guide our understanding of how the upper convex hull reacts to to small changes in the network.

As a final remark on the first track for future work, another interesting topic is the size of the upper convex hulls of P and N . While these two sets typically grow exponentially in depth, it is only $\mathcal{U}^*(P)$ and $\mathcal{U}^*(N)$ that matter. Conjecture 5.1.20 predicts that usually $|\mathcal{U}^*(P)| \ll |P|$, $|\mathcal{U}^*(N)| \ll |N|$. Better understanding the size of the upper convex hulls would also facilitate our understanding of the volume hypothesis.

This brings us to the second track. While we have gathered evidence for, and provided a framework to think about, the volume hypothesis, we still lack a general theorem which provides e.g. bounds for the distribution of the number of affine regions, given a fixed architecture. For future work, it would be interesting to derive more general results which rigorously establish the volume hypothesis.

Bibliography

- [1] P. Nakkiran, G. Kaplun, Y. Bansal, T. Yang, B. Barak, and I. Sutskever, “Deep Double Descent: Where Bigger Models and More Data Hurt,” *Journal of Statistical Mechanics: Theory and Experiment*, vol. 2021, no. 12, p. 124003, 2021.
- [2] C. Zhang, S. Bengio, M. Hardt, B. Recht, and O. Vinyals, “Understanding Deep Learning (Still) Requires Rethinking Generalization,” *Communications of the ACM*, vol. 64, no. 3, pp. 107–115, 2021.
- [3] D. Soudry, E. Hoffer, M. S. Nacson, S. Gunasekar, and N. Srebro, “The Implicit Bias of Gradient Descent on Separable Data,” *Journal of Machine Learning Research*, vol. 19, no. 70, pp. 1–57, 2018.
- [4] C. Zhang, Q. Liao, A. Rakhlin, K. Sridharan, B. Miranda, N. Golowich, and T. Poggio, “Musings on Deep Learning: Properties of SGD,” Center for Brains, Minds and Machines (CBMM), Tech. Rep., 2017.
- [5] Y. Zhang, A. M. Saxe, M. S. Advani, and A. A. Lee, “Energy–Entropy Competition and the Effectiveness of Stochastic Gradient Descent in Machine Learning,” *Molecular Physics*, vol. 116, no. 21–22, pp. 3214–3223, 2018.
- [6] P.-y. Chiang, R. Ni, D. Y. Miller, A. Bansal, J. Geiping, M. Goldblum, and T. Goldstein, “Loss Landscapes are All You Need: Neural Network Generalization Can Be Explained Without the Implicit Bias of Gradient Descent,” in *The Eleventh International Conference on Learning Representations*, 2022.
- [7] G. Valle-Perez, C. Q. Camargo, and A. A. Louis, “Deep Learning Generalizes Because the Parameter-Function Map is Biased Towards Simple Functions,” *arXiv preprint arXiv:1805.08522*, 2018.
- [8] K. Dingle, C. Q. Camargo, and A. A. Louis, “Input–Output Maps Are Strongly Biased Towards Simple Outputs,” *Nature communications*, vol. 9, no. 1, p. 761, 2018.
- [9] C. Mingard, H. Rees, G. Valle-Pérez, and A. A. Louis, “Do Deep Neural Networks Have an Inbuilt Occam’s Razor?” *arXiv preprint arXiv:2304.06670*, 2023.
- [10] C. Zhang, S. Bengio, M. Hardt, B. Recht, and O. Vinyals, “Understanding Deep Learning Requires Rethinking Generalization,” *arXiv preprint arXiv:1611.03530*, 2016.
- [11] G. F. Montufar, R. Pascanu, K. Cho, and Y. Bengio, “On the Number of Linear

- Regions of Deep Neural Networks," *Advances in neural information processing systems*, vol. 27, 2014.
- [12] L. Zhang, G. Naitzat, and L.-H. Lim, "Tropical Geometry of Deep Neural Networks," in *International Conference on Machine Learning*. PMLR, 2018, pp. 5824–5832.
- [13] M. Raghu, B. Poole, J. Kleinberg, S. Ganguli, and J. Sohl-Dickstein, "On the Expressive Power of Deep Neural Networks," in *international conference on machine learning*. PMLR, 2017, pp. 2847–2854.
- [14] B. Hanin and D. Rolnick, "Deep Relu Networks Have Surprisingly Few Activation Patterns," *Advances in neural information processing systems*, vol. 32, 2019.
- [15] M. Telgarsky, "Representation Benefits of Deep Feedforward Networks," *arXiv preprint arXiv:1509.08101*, 2015.
- [16] B. Hanin and D. Rolnick, "Complexity of Linear Regions in Deep Networks," in *International Conference on Machine Learning*. PMLR, 2019, pp. 2596–2604.
- [17] R. Pascanu, G. Montufar, and Y. Bengio, "On the Number of Response Regions of Deep Feed Forward Networks with Piece-Wise Linear Activations," *arXiv preprint arXiv:1312.6098*, 2013.
- [18] P. Piwek, A. Klukowski, and T. Hu, "Exact Count of Boundary Pieces of Relu Classifiers: Towards the Proper Complexity Measure for Classification," in *Uncertainty in Artificial Intelligence*. PMLR, 2023, pp. 1673–1683.
- [19] A. Goujon, A. Etemadi, and M. Unser, "On the Number of Regions of Piecewise Linear Neural Networks," *Journal of Computational and Applied Mathematics*, vol. 441, p. 115667, 2024.
- [20] C. Keup and M. Helias, "Origami in N Dimensions: How Feed-Forward Networks Manufacture Linear Separability," *arXiv preprint arXiv:2203.11355*, 2022.
- [21] M. Lewandowski, H. Eghbalzadeh, B. Heinzl, R. Pisoni, and B. A. Moser, "On Space Folds of Relu Neural Networks," *arXiv preprint arXiv:2502.09954*, 2025.
- [22] R. Balestrieri, A. I. Humayun, and R. Baraniuk, "On the Geometry of Deep Learning," *arXiv preprint arXiv:2408.04809*, 2024.
- [23] M. Alfarra, A. Bibi, H. Hammoud, M. Gaafar, and B. Ghanem, "On the Decision Boundaries of Neural Networks: A Tropical Geometry Perspective," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 45, no. 4, pp. 5027–5037, 2022.
- [24] C. Nwankpa, W. Ijomah, A. Gachagan, and S. Marshall, "Activation Functions: Comparison of Trends in Practice and Research for Deep Learning," *arXiv preprint arXiv:1811.03378*, 2018.
- [25] S. Gu, B. Kelly, and D. Xiu, "Empirical Asset Pricing via Machine Learning," *The Review of Financial Studies*, vol. 33, no. 5, pp. 2223–2273, 2020.

- [26] Y. LeCun, L. Bottou, Y. Bengio, and P. Haffner, "Gradient-Based Learning Applied to Document Recognition," *Proceedings of the IEEE*, vol. 86, no. 11, pp. 2278–2324, 1998.
- [27] A. Krizhevsky, I. Sutskever, and G. E. Hinton, "Imagenet Classification with Deep Convolutional Neural Networks," *Advances in neural information processing systems*, vol. 25, 2012.
- [28] D. P. Kingma and J. Ba, "Adam: A Method for Stochastic Optimization," in *International Conference on Learning Representations*, 2015, pp. 1–15.
- [29] D. Maclagan and B. Sturmfels, *Introduction to Tropical Geometry*. American Mathematical Society, 2021, vol. 161.
- [30] K. Chandrasekaran, "Lecture Notes on Linear Programming. Lecture 5."
- [31] G. (https://math.stackexchange.com/users/83800/gary), "Sum of Reciprocal Sine Function," Mathematics Stack Exchange, uRL:https://math.stackexchange.com/q/4223978 (version: 2024-04-17). [Online]. Available: https://math.stackexchange.com/q/4223978
- [32] S. Koumandos, "Remarks on some Completely Monotonic Functions," *Journal of Mathematical Analysis and Applications*, vol. 324, no. 2, pp. 1458–1461, 2006.
- [33] I. S. Gradshteyn and I. M. Ryzhik, *Table of Integrals, Series, and Products*. Academic press, 2014.
- [34] R. L. Graham, *Concrete Mathematics: A Foundation for Computer Science*. Pearson Education India, 1994.
- [35] A. Rényi and R. Sulanke, "Über die konvexe hülle von n zufällig gewählten punkten," *Zeitschrift für Wahrscheinlichkeitstheorie und verwandte Gebiete*, vol. 2, pp. 75–84, 1963.
- [36] L. A. Santaló, *Integral Geometry and Geometric Probability*. Cambridge university press, 2004.
- [37] M. P. Knapp, "Sines and Cosines of Angles in Arithmetic Progression," *Mathematics magazine*, vol. 82, no. 5, pp. 371–372, 2009.

Appendix A

Helping Lemmas

This appendix contains a number of helpful lemmas.

Lemma A.0.1. *Let $X \sim \mathcal{N}(0, \sigma_x^2)$ and $Y \sim \mathcal{N}(0, \sigma_y^2)$ be independent Gaussian random variables. Then*

$$\mathbb{P}(0 \leq X \leq Y) = \mathbb{P}(0 \geq X \geq Y) = \frac{1}{2\pi} \arctan\left(\frac{\sigma_y}{\sigma_x}\right).$$

Proof. Note that (X, Y) is a multivariate Gaussian with density

$$p_{(X,Y)}(x, y) = \frac{1}{2\pi\sigma_x\sigma_y} e^{-\left(\frac{x^2}{2\sigma_x^2} + \frac{y^2}{2\sigma_y^2}\right)}.$$

Apply a change of variables like $x = \sqrt{2}r\sigma_x \cos \theta$, $y = \sqrt{2}r\sigma_y \sin \theta$, which has Jacobian $|J| = 2r\sigma_x\sigma_y$ and

$$y \geq x \geq 0 \iff \frac{\sigma_x}{\sigma_y} \leq \tan \theta.$$

We can thus compute

$$\begin{aligned} \mathbb{P}(0 < X < Y) &= \frac{1}{2\pi\sigma_x\sigma_y} \int_{\substack{(x,y) \in \mathbb{R}^2 \\ 0 < x < y}} e^{-\left(\frac{x^2}{2\sigma_x^2} + \frac{y^2}{2\sigma_y^2}\right)} dx dy \\ &= \frac{1}{\pi} \int_{\arctan\left(\frac{\sigma_x}{\sigma_y}\right)}^{\pi/2} d\theta \int_0^\infty dr e^{-r^2} \\ &= \frac{1}{4} - \frac{1}{2\pi} \arctan\left(\frac{\sigma_x}{\sigma_y}\right) = \frac{1}{2\pi} \arctan\left(\frac{\sigma_y}{\sigma_x}\right), \end{aligned}$$

where in the last step we used the identity $\arctan\left(\frac{1}{x}\right) = \frac{\pi}{2} - \arctan(x)$. □

The following three Lemmas contain well known trigonometric identities:

Lemma A.0.2. *Given two angles $\alpha, \beta \in \mathbb{R}$, the following identities hold:*

$$\begin{aligned}\sin \alpha + \sin \beta &= 2 \sin \left(\frac{\alpha + \beta}{2} \right) \cos \left(\frac{\alpha - \beta}{2} \right) \\ \cos \alpha + \cos \beta &= 2 \cos \left(\frac{\alpha + \beta}{2} \right) \cos \left(\frac{\alpha - \beta}{2} \right).\end{aligned}$$

Lemma A.0.3. *Given two angles $\alpha, \beta \in \mathbb{R}$, the following identities hold:*

$$\begin{aligned}\cos \alpha - \cos \beta &= -2 \sin \left(\frac{\alpha + \beta}{2} \right) \sin \left(\frac{\alpha - \beta}{2} \right) \\ \sin \alpha - \sin \beta &= 2 \cos \left(\frac{\alpha + \beta}{2} \right) \sin \left(\frac{\alpha - \beta}{2} \right).\end{aligned}$$

Lemma A.0.4. *Given two angles $\alpha, \beta \in \mathbb{R}$, the following identity holds:*

$$\sin(\alpha - \beta) = \sin \alpha \cos \beta - \cos \alpha \sin \beta.$$

Lemma A.0.5. *Given two angles $\alpha, \beta \in \mathbb{R}$, the following identity holds:*

$$\sin(\alpha) \cos(\beta) = \frac{1}{2} [\sin(\alpha + \beta) + \sin(\alpha - \beta)]$$

Lemma A.0.6 ([37, Theorem 1]). *For any $a, d \in \mathbb{R}$, the following is true:*

$$\sum_{k=0}^{n-1} \sin(a + k \cdot d) = \frac{\sin(n \cdot \frac{d}{2})}{\sin(\frac{d}{2})} \cdot \sin \left(a + (n-1) \frac{d}{2} \right).$$