The Separation Capacity of Random Neural Networks

Sjoerd Dirksen[*]   **Martin Genzel**[*]   **Laurent Jacques**[†]   **Alexander Stollenwerk**[†]

Abstract. Neural networks with random weights appear in a variety of machine learning applications, most prominently as the initialization of many deep learning algorithms and as a computationally cheap alternative to fully learned neural networks. In the present article we enhance the theoretical understanding of random neural nets by addressing the following data separation problem: under what conditions can a random neural network make two classes $\mathcal{X}^-, \mathcal{X}^+ \subset \mathbb{R}^d$ (with positive distance) linearly separable? We show that a sufficiently large two-layer ReLU-network with standard Gaussian weights and uniformly distributed biases can solve this problem with high probability. Crucially, the number of required neurons is explicitly linked to geometric properties of the underlying sets $\mathcal{X}^-, \mathcal{X}^+$ and their mutual arrangement. This instance-specific viewpoint allows us to overcome the usual curse of dimensionality (exponential width of the layers) in non-pathological situations where the data carries low-complexity structure. We quantify the relevant structure of the data in terms of a novel notion of mutual complexity (based on a localized version of Gaussian mean width), which leads to sound and informative separation guarantees. We connect our result with related lines of work on approximation, memorization, and generalization.

Key words. Random neural networks, classification, hyperplane separation, high-dimensional geometry, Gaussian mean width.

1 Introduction

Despite the unprecedented success of neural networks (NNs) in countless applications [LBH15; Sch15; GBC16], a rigorous understanding of their operating principles is still in its infancy. The present work is devoted to a mathematical study of *random NNs*, i.e., feedforward NNs whose weight parameters are drawn from a generic probability distribution. Random NNs play an important role in machine learning in at least three different ways. First, it is standard to initialize the training of a (deep) NN using a random NN and it is well known that this initialization is a key contributor to the exceptional performance of NNs [HZRS15; GBC16; AB19]. Second, it has been empirically observed that architecture search can be effectively carried out with random NNs, in the sense that the hierarchy in performance of fully trained architectures closely matches the hierarchy of the architectures with random weights [Sax+11]. Finally, random NNs have been extensively investigated as a cheap computational alternative to fully trained neural networks: it has been observed empirically that pre-processing with a random NN and applying a simple classification method already gives surprisingly good results [HZS06; RR08; ZBHRV17]. For these reasons, it is of substantial interest to gain a deeper theoretical understanding of the properties of random neural networks.
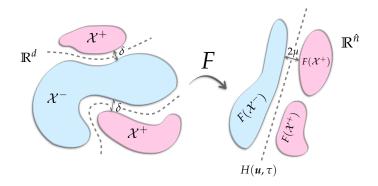

Figure 1: **Illustration of Problem 1.1.** Can a random NN $F\colon \mathbb{R}^d \to \mathbb{R}^{\hat{n}}$ "disentangle" the two sets $\mathcal{X}^-, \mathcal{X}^+ \subset \mathbb{R}^d$ such that they become linearly separable in the feature space $\mathbb{R}^{\hat{n}}$ with a positive margin μ? Except for being δ-separated and bounded, $\mathcal{X}^-$ and $\mathcal{X}^+$ may have an arbitrary decision boundary, possibly with multiple connected components.

---

[*]Utrecht University, Mathematical Institute, Utrecht, Netherlands (e-mails: [s.dirksen, m.genzel]@uu.nl).
[†]ISPGroup, INMA, ICTEAM Institute, UCLouvain, Louvain-la-Neuve, Belgium (e-mails: [laurent.jacques, alexander.stollenwerk]@uclouvain.be).

In this work, we shed new light on the capabilities of random neural networks as a pre-processor by addressing the following fundamental problem on class separability:

> **Problem 1.1** Consider two bounded, possibly infinite sets $\mathcal{X}^-, \mathcal{X}^+ \subset \mathbb{R}^d$ that are *δ-separated*, i.e.,
>
> $$\|x^+ - x^-\|_2 \geq \delta \qquad \text{for all } x^+ \in \mathcal{X}^+ \text{ and } x^- \in \mathcal{X}^-.$$
>
> Let $F\colon \mathbb{R}^d \to \mathbb{R}^{\hat{n}}$ represent a (multi-layer) feedforward NN with *random weights*, where the architecture of F may depend on $\mathcal{X}^-$ and $\mathcal{X}^+$.
>
> *Under what conditions does F make the classes $\mathcal{X}^-$ and $\mathcal{X}^+$ linearly separable with high probability? Is there a lower bound for the induced separation margin?*

Formally, we will identify conditions that ensure the existence (with high probability) of a hyperplane $H[u, \tau] := \{z \in \mathbb{R}^{\hat{n}} \mid \langle u, z \rangle + \tau = 0\}$ with $\|u\|_2 = 1$ and $\tau \in \mathbb{R}$ that separates $F(\mathcal{X}^-)$ and $F(\mathcal{X}^+)$ with a certain *margin* $\mu > 0$, i.e.,

$$
\begin{aligned}
\langle u, F(x^-) \rangle + \tau &\leq -\mu \qquad \text{for all } x^- \in \mathcal{X}^-, \\
\langle u, F(x^+) \rangle + \tau &\geq +\mu \qquad \text{for all } x^+ \in \mathcal{X}^+.
\end{aligned}
\tag{1.1}
$$

Problem 1.1 thus states a purely geometric question on the *separation capacity* of random NNs, see Figure 1 for an illustration. However, it is useful to bear in mind that the ability to render two "intertwined" sets linearly separable also has immediate consequences for associated learning tasks. To see this, let (X, Y) be an arbitrary data distribution on $\mathbb{R}^d \times \{\pm 1\}$ with $\mathbb{P}(X \in \mathcal{X}^+ \mid Y = +1) = 1 = \mathbb{P}(X \in \mathcal{X}^- \mid Y = -1)$, i.e., the binary label Y is consistent with the classes $\mathcal{X}^-$ and $\mathcal{X}^+$. Conditioned on the high-probability event of Problem 1.1, the transformed pair $(F(X), Y)$ then fulfills a *hard-margin condition*:

$$\mathbb{P}(Y \cdot (\langle u, F(X) \rangle + \tau) \geq \mu) = 1, \tag{1.2}$$

where $H[u, \tau]$ denotes the separating hyperplane in (1.1). Given i.i.d. training samples of (X, Y), this enables us to learn the unknown output parameters (u, τ) by standard classification methods, such as support vector machines [SC08]. In particular, one can achieve provable control over the generalization error in terms of the margin size μ, see e.g. [SB14, Thm. 15.4]. Of similar relevance is the width $\hat{n}$ of the output-layer of F, as it determines the ambient dimension of the feature space and hence the computational complexity of the classification method. For these reasons, we seek to solve Problem 1.1 with reasonable bounds for both μ and $\hat{n}$.

In principle, the aforementioned idea of using random NNs as a pre-processing step for well-understood (linear) classifiers is not new [HZS06; RR08; ZBHRV17]. But despite conceptual overlaps, the analytical approach of the present article is different from most existing works, see also Section 1.5 for a short literature overview. Although Problem 1.1 includes a large family of classification tasks — namely all pairs of δ-separated sets — we are primarily interested in an *instance-specific* analysis: our main results quantify the dependence of the key parameters μ and $\hat{n}$ as functions of the underlying classes $\mathcal{X}^-$ and $\mathcal{X}^+$ and their "interaction". The resulting instance-specific bounds allow us to avoid pessimistic (worst-case) bounds caused by pathological pairs of sets. In our analysis we make no explicit assumptions about the data in Problem 1.1, such as a handcrafted generative model or sampling from a generic distribution. Instead, we will introduce complexity measures that quantify the geometric complexities of $\mathcal{X}^-$ and $\mathcal{X}^+$ as well as their mutual entanglement. This perspective appears more natural to us in the context of data-driven methods.

Let us now specify the class of random NNs for which we will explore Problem 1.1. Throughout, the function $F\colon \mathbb{R}^d \to \mathbb{R}^{\hat{n}}$ will be composed of (hidden) layers of the following form.

Definition 1.2 We call $\Phi\colon \mathbb{R}^{n_{\text{in}}} \to \mathbb{R}^{n_{\text{out}}}$ a *random ReLU-layer* with maximal bias $\lambda \geq 0$ if

$$\Phi(x) = \sqrt{\tfrac{2}{n_{\text{out}}}} \cdot \text{ReLU}(Wx + b), \qquad x \in \mathbb{R}^{n_{\text{in}}}, \tag{1.3}$$

where the weight matrix $W \in \mathbb{R}^{n_{\text{out}} \times n_{\text{in}}}$ has standard Gaussian entries, the bias vector b is uniformly distributed on $[-\lambda, \lambda]^{n_{\text{out}}}$, independently of W, and the element-wise activation function is the rectified linear unit (ReLU), i.e., $\text{ReLU}(s) := \max\{0, s\}$ for $s \in \mathbb{R}$.

It will turn out that already two random ReLU-layers are sufficient for our solution to Problem 1.1, although in principle deeper architectures are also possible. We consider the ReLU mainly because of its popularity, but our analysis can be adapted for other common activation functions, e.g., the thresholding activation. Let us note that the random weights and normalization in Definition 1.2 do not correspond to a standard initialization in deep learning. The closest is the popular *He initialization* [HZRS15], which would correspond to replacing n_{out} by n_{in} and taking $\boldsymbol{b} = \boldsymbol{0}$ in (1.3). The reason for our non-standard choice will become clear in our proof sketch in Section 1.4.

Instead of directly formulating our main result, Theorem 2.3, we will first present several accessible special cases of increasing generality. Afterwards, we will highlight our proof strategy in Section 1.4.

1.1 A Gentle Start: Finite Sets and Memorization

Our first result below gives an answer to Problem 1.1 in the situation of *finite* point sets. In the following, $\mathbb{B}_2^d := \{\boldsymbol{x} \in \mathbb{R}^d \mid \|\boldsymbol{x}\|_2 \leq 1\}$ denotes the Euclidean unit ball; see also Section 1.6 for a summary of common notation used in this article.

Theorem 1.3 (Finite sets) *There exist absolute constants $c, C > 0$ such that the following holds. Let $\mathcal{X}^-, \mathcal{X}^+ \subset \mathbb{B}_2^d$ be δ-separated sets with $N^+ := |\mathcal{X}^+| < \infty$ and $N^- := |\mathcal{X}^-| < \infty$. Suppose that $\lambda \gtrsim \sqrt{\log(e\lambda/\delta)}$. Let $\Phi \colon \mathbb{R}^d \to \mathbb{R}^n$ and $\hat{\Phi} \colon \mathbb{R}^n \to \mathbb{R}^{\hat{n}}$ be two (independent) random ReLU-layers with maximal biases $\lambda, \hat{\lambda} \geq 0$, respectively, such that*

$$n \gtrsim \left(\tfrac{\lambda}{\delta}\right)^8 \cdot \log(2N^- N^+/\eta) \tag{1.4}$$

and

$$\hat{\lambda} \gtrsim \left(\tfrac{\lambda}{\delta}\right)^4 \cdot \left(\alpha + \lambda\right), \qquad \hat{n} \gtrsim \tfrac{\hat{\lambda}}{\lambda} \cdot \theta \cdot \log(N^-/\eta), \tag{1.5}$$

where $\alpha = \sqrt{\log N^+}$ and

$$\theta = \exp\left(C \cdot \left(\alpha^2 + \lambda^2\right) \cdot \lambda^6 \cdot \delta^{-8} \cdot \log(\lambda/\delta)\right). \tag{1.6}$$

Then, given the two-layer random neural network $F \colon \mathbb{R}^d \to \mathbb{R}^{\hat{n}}$, $\boldsymbol{x} \mapsto \hat{\Phi}(\Phi(\boldsymbol{x}))$, with probability at least $1 - \eta$, the sets $F(\mathcal{X}^-), F(\mathcal{X}^+) \subset \hat{\lambda}\mathbb{B}_2^{\hat{n}}$ are linearly separable with margin $c\lambda^2/(\hat{\lambda}\theta)$.

The governing condition in this result (and in all following results, Theorems 1.4, 1.5, and 2.3 below) is condition (1.5) on the width $\hat{n}$ of the second layer. It features the term θ that scales exponentially in terms of the logarithm of the number of points, so that $\hat{n}$ scales as $N_+^{\text{poly}(\lambda, 1/\delta)}$, in contrast to the logarithmic scaling of n in (1.4). To gauge whether this condition is necessary, let us connect Theorem 1.3 to the *memorization capacity* of random NNs. The ability of memorizing large data sets (including their noisy components) is a well-known phenomenon in deep learning research and considered as an important piece of the still unsolved generalization puzzle [ZBHRV17; ZBHRV21]. Theorem 1.3 applies to any (δ-separated) completely unstructured dataset — imagine a point cloud with arbitrary binary labels. Remarkably, one can therefore memorize the labels of any such set with high probability by efficiently computing a separating hyperplane of $F(\mathcal{X}^-)$ and $F(\mathcal{X}^+)$, e.g., using a hard-margin SVM.[1]

Although this shows that random NNs can be powerful memorizers in practice, existing results in the literature indicate that perfect memorization is already possible when the number of neurons scales linearly in $(N^- + N^+)$ up to logarithmic factors, see e.g. [YSJ19; BN20; Ver20]. Hence, we expect that the dependence on δ and λ within the exponential term θ in (1.5) may be improved.

1.2 Separation of Euclidean Balls

Although Theorem 1.3 provides a margin bound, its actual size was not relevant to the network's memorization capacity. The situation is different for *infinite classes*, on which we will focus from

---

[1]This observation is especially interesting when no obvious learning rule is available. On the other hand, if the data carries more structure (e.g., $\mathcal{X}^-$ and $\mathcal{X}^+$ are already linearly separable), there certainly exist more effective approaches than randomized transforms.
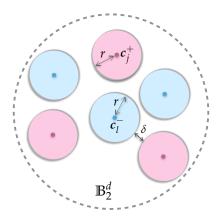
Figure 2: **Illustration of the Euclidean-ball model in Theorem 1.4.** The sets $\mathcal{X}^-, \mathcal{X}^+ \subset \mathbb{B}_2^d$ consist of Euclidean balls of radius r, where the center points are denoted by c_l^- and c_j^+, respectively. Note that the δ-separation only concerns balls of different classes, while arbitrary intersections are allowed within each class.

now on. Problem 1.1 is then connected to a binary classification task through the hard-margin condition (1.2), and the margin size determines the generalization performance [SB14, Thm. 15.4]. Our next result may be seen as a natural extension of Theorem 1.3, replacing discrete data points by a finite collection of Euclidean balls; see Figure 2 for an illustration of this model.

Theorem 1.4 (Euclidean balls) *There exist absolute constants $c, C > 0$ such that the following holds.*
Let $\mathcal{X}^-, \mathcal{X}^+ \subset \mathbb{B}_2^d$ be δ-separated sets that can be written as the union of finitely many Euclidean balls of radius $r \geq 0$, i.e.,

$$\mathcal{X}^- = \bigcup_{l \in [N^-]} \mathbb{B}_2^d(c_l^-, r), \quad \mathcal{X}^+ = \bigcup_{j \in [N^+]} \mathbb{B}_2^d(c_j^+, r).$$

Suppose that $\lambda \gtrsim \sqrt{\log(e\lambda/\delta)}$ and $r \lesssim \delta^2/\lambda$. We assume that $\Phi \colon \mathbb{R}^d \to \mathbb{R}^n$ and $\hat{\Phi} \colon \mathbb{R}^n \to \mathbb{R}^{\hat{n}}$ are two (independent) random ReLU-layers with maximal biases $\lambda, \hat{\lambda} \geq 0$, respectively, such that

$$n \gtrsim (1 + \lambda^6 \delta^{-8} r^2) \cdot d + \left(\tfrac{\lambda}{\delta}\right)^8 \cdot \log(2N^- N^+/\eta) \tag{1.7}$$

and (1.5) holds with $\alpha = r\sqrt{d} + \sqrt{\log N^+}$ and θ as in (1.6). Then, given the two-layer random neural network $F \colon \mathbb{R}^d \to \mathbb{R}^{\hat{n}}$, $x \mapsto \hat{\Phi}(\Phi(x))$, with probability at least $1 - \eta$, the sets $F(\mathcal{X}^-), F(\mathcal{X}^+) \subset \hat{\lambda}\mathbb{B}_2^{\hat{n}}$ are linearly separable with margin $c\lambda^2/(\hat{\lambda}\theta)$.

It is worth noting that in the limit case $r = 0$, the above statement is essentially consistent with Theorem 1.3. On the other hand, Theorem 1.4 reveals the price of dealing with full-dimensional sets instead of points: the bound on the output dimension $\hat{n}$ scales exponentially in terms of $r^2 d$. Thus, to avoid the curse of dimensionality, the radius needs to satisfy $r \lesssim 1/\sqrt{d}$. Under this assumption, Theorem 1.4 certifies that random NNs can efficiently separate (unstructured) collections of Euclidean balls.

1.3 Towards a General Separation Guarantee

So far, we have only considered specific examples of data sets. Our next theorem concerns Problem 1.1 for arbitrary δ-separated classes. For a formal statement, we need to introduce two important geometric parameters. The *covering number* of a bounded subset $\mathcal{X} \subset \mathbb{R}^d$ at scale $r > 0$ is given by

$$\mathcal{N}(\mathcal{X}, r) := \min \left\{ N \in \mathbb{N} \mid \exists c_1, \ldots, c_N \in \mathbb{R}^d \colon \mathcal{X} \subset \bigcup_{j \in [N]} \mathbb{B}_2^d(c_j, r) \right\}, \tag{1.8}$$

i.e., the smallest number of Euclidean balls of radius r required to cover $\mathcal{X}$. Moreover, the *(Gaussian) mean width* of $\mathcal{X}$ is defined as

$$w(\mathcal{X}) := \mathbb{E}_g \left[\sup_{x \in \mathcal{X}} \langle g, x \rangle \right],$$

where $g \sim \mathsf{N}(\mathbf{0}, I_d)$ denotes a standard Gaussian random vector. Both $\mathcal{N}(\mathcal{X}, r)$ and $w(\mathcal{X})$ are natural complexity measures, which are well-established in high-dimensional geometry, statistics, and signal processing, see e.g. [GM04; CRPW12; Tal14; Ver18].

Theorem 1.5 (General sets) *There exist absolute constants $c, c', C > 0$ such that the following holds. Let $\mathcal{X}^-, \mathcal{X}^+ \subset \mathbb{B}_2^d$ be δ-separated sets and let $\lambda \gtrsim \sqrt{\log(e\lambda/\delta)}$. We set $N^- := \mathcal{N}(\mathcal{X}^-, c\delta^2/\lambda)$ and $N^+ := \mathcal{N}(\mathcal{X}^+, c\delta^2/\lambda)$. We assume that $\Phi \colon \mathbb{R}^d \to \mathbb{R}^n$ and $\hat{\Phi} \colon \mathbb{R}^n \to \mathbb{R}^{\hat{n}}$ are two (independent) random ReLU-layers with maximal biases $\lambda, \hat{\lambda} \geq 0$, respectively, such that[2]*

$$n \gtrsim w^2(\mathrm{cone}(\mathcal{X}^- - \mathcal{X}^-) \cap \mathsf{S}^{d-1}) + w^2(\mathrm{cone}(\mathcal{X}^+ - \mathcal{X}^+) \cap \mathsf{S}^{d-1}),$$
$$n \gtrsim \left(\tfrac{\lambda}{\delta}\right)^8 \cdot \left(\lambda^{-2}\left(w^2(\mathcal{X}^-) + w^2(\mathcal{X}^+)\right) + \log(2N^- N^+/\eta)\right) \tag{1.9}$$

and (1.5) holds with $\alpha = w(\mathcal{X}^-) + w(\mathcal{X}^+)$ and θ as in (1.6). Then, given the two-layer random neural network $F \colon \mathbb{R}^d \to \mathbb{R}^{\hat{n}}, \; x \mapsto \hat{\Phi}(\Phi(x))$, with probability at least $1 - \eta$, the sets $F(\mathcal{X}^-), F(\mathcal{X}^+) \subset \hat{\lambda}\mathbb{B}_2^{\hat{n}}$ are linearly separable with margin $c\lambda^2/(\hat{\lambda}\theta)$.

Compared to Theorem 1.4, Theorem 1.5 yields a much stronger statement due to the use of the mean width as a complexity measure. To see this, let $\mathcal{X}^+ \subset \bigcup_{j \in [N^+]} \mathbb{B}_2^d(c_j^+, r)$ be any covering of $\mathcal{X}^+$ at scale $r = c\lambda^{-1}\delta^2$ and consider the following upper bound (see Lemma 6.1):

$$w(\mathcal{X}^+) \lesssim w^+ + \sqrt{\log N^+}, \tag{1.10}$$

where $w^+ = \max_{j \in [N^+]} w(\mathcal{X}^+ \cap \mathbb{B}_2^d(c_j^+, r))$. An analogous bound holds for $\mathcal{X}^-$. While using the worst-case estimate $w^+ \lesssim r\sqrt{d}$ would lead to a similar bottleneck as in Theorem 1.4, the *localized* mean width parameter w^+ can be substantially smaller for structured data sets. Typical examples are data residing on a low-dimensional manifold or contained in the convex hull of finitely many points. On the other hand, the covering number N^+ reflects the *global* size of $\mathcal{X}^+$ in (1.10).

For these reasons, Theorem 1.5 takes an important step towards a general solution to Problem 1.1, which meets our overall goal of *instance-specific* bounds for the separation margin and layer widths. Having said this, the geometric parameters in this result only capture the individual complexities of $\mathcal{X}^-$ and $\mathcal{X}^+$, but remain silent about their mutual arrangement. For instance, one would expect that two sets become easier to separate if their "centers of mass" are farther apart, even though the minimal distance δ is small; see Figure 5 in Section 2 for an illustration. In such scenarios, a non-uniform covering strategy for $\mathcal{X}^-$ and $\mathcal{X}^+$ is preferable, in the sense that data points far away from the decision boundary should be covered by fewer but larger balls. The most general outcome of this work, Theorem 2.3, makes this intuition precise by employing a novel notion of *mutual complexity* (see Definition 2.1 and 2.2). We refer to Section 2 for an in-depth discussion and further refinements due to Theorem 2.3. Finally, we emphasize that all previously presented results follow from Theorem 2.3 as special cases, see Section 6 for detailed proofs.

1.4 Proof Strategy

To keep our exposition as simple as possible, we will describe our proof strategy in the prototypical situation of Euclidean balls from Theorem 1.4, see also Figure 2. Recall that $\mathcal{C}^- := \{c_1^-, \ldots, c_{N^-}^-\}$ and $\mathcal{C}^+ := \{c_1^+, \ldots, c_{N^+}^+\}$ denote the center points of $\mathcal{X}^-$ and $\mathcal{X}^+$, respectively. For convenience, we also set $\mathcal{X}_l^- := \mathbb{B}_2^d(c_l^-, r)$ and $\mathcal{X}_j^+ := \mathbb{B}_2^d(c_j^+, r)$ so that $\mathcal{X}^- = \bigcup_{l \in [N^-]} \mathcal{X}_l^-$ and $\mathcal{X}^+ = \bigcup_{j \in [N^+]} \mathcal{X}_j^+$.

Our data separation approach consists of a *two-step procedure*, which essentially corresponds to the composition of the random layers $\Phi \colon \mathbb{R}^d \to \mathbb{R}^n$ and $\hat{\Phi} \colon \mathbb{R}^n \to \mathbb{R}^{\hat{n}}$. Although Φ and $\hat{\Phi}$ both follow exactly the same random design (see Definition 1.2), we will see that their purposes are different: while the first one already establishes a desirable geometrical configuration under mild conditions, a major challenge is to show that one can actually take advantage of it by applying a second (wider) random layer.

---

[2]Note that the first line in (1.9) is always satisfied if $n \gtrsim d$, since we have that $w^2(\mathsf{S}^{d-1}) \asymp d$. See also Section 1.6 for a precise definition of $\mathrm{cone}(\cdot)$.
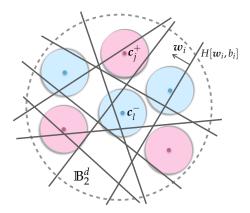
Figure 3: **Random hyperplanes in the input domain.** Each coordinate of the first layer $\Phi \colon \mathbb{R}^d \to \mathbb{R}^n$ can be associated with a random hyperplane $H[\boldsymbol{w}_i, b_i]$, where $\boldsymbol{w}_i \sim \mathsf{N}(\boldsymbol{0}, \boldsymbol{I}_d)$ and b_i is uniformly distributed on $[-\lambda, \lambda]$. Already a relatively few of such hyperplanes are enough to separate every pair of center points $(\boldsymbol{c}_l^-, \boldsymbol{c}_j^+) \in \mathcal{C}^- \times \mathcal{C}^+$ at least once. For this, a sufficiently large bias parameter ($\lambda \gtrsim 1$) is vital, as it ensures a *uniform tessellation* of the input domain $\mathbb{B}_2^d$; otherwise, the probability of separating points close to the boundary of $\mathbb{B}_2^d$ would become too low.

To understand the effect of the first layer Φ, it is useful to take a coordinate-wise perspective:

$$[\Phi(\boldsymbol{x})]_i = \sqrt{\tfrac{2}{n}}\, \mathrm{ReLU}(\langle \boldsymbol{w}_i, \boldsymbol{x} \rangle + b_i), \quad i = 1, \ldots, n, \quad \boldsymbol{x} \in \mathbb{R}^d,$$

where $\boldsymbol{w}_i \in \mathbb{R}^d$ is the i-th row of the weight matrix $\boldsymbol{W} \in \mathbb{R}^{n \times d}$ and $b_i \in [-\lambda, \lambda]$ the corresponding bias. Thus, $[\Phi(\boldsymbol{x})]_i$ indicates on which side of the (unnormalized) random hyperplane $H[\boldsymbol{w}_i, b_i]$ a given point $\boldsymbol{x} \in \mathbb{R}^d$ lies. Our first major proof step (elaborated in Theorem 5.3) shows that as long as $\lambda \gtrsim 1$, the following holds with high probability: for every pair of center points $(\boldsymbol{c}_l^-, \boldsymbol{c}_j^+) \in \mathcal{C}^- \times \mathcal{C}^+$, there are coordinates $I_{l,j} \subset [n]$ with $|I_{l,j}| \gtrsim \delta \lambda^{-1} n$ such that $H[\boldsymbol{w}_i, b_i]$ separates $\boldsymbol{c}_l^-$ from $\boldsymbol{c}_j^+$ for all $i \in I_{l,j}$, in fact

$$[\Phi(\boldsymbol{c}_l^-)]_i = 0 \quad \text{and} \quad [\Phi(\boldsymbol{c}_j^+)]_i \gtrsim \tfrac{\delta}{\sqrt{n}}, \tag{1.11}$$

see Figure 3 for an illustration. Here, it is crucial to note that the corresponding coordinates in $\Phi(\boldsymbol{c}_l^-)$ are actually vanishing, due to the *non-linear* activation. The key insight to show (1.11) is that the probability of a single random hyperplane separating a fixed pair of δ-separated points is of order $\Omega(\delta \lambda^{-1})$, see Theorem 3.7. Combining this with a Chernoff bound over all hyperplanes associated with Φ then leads to a high-probability event of the above type. A remarkable fact about this argument is that the required layer width n is very moderate, scaling only logarithmically with the number of center points (see (1.7)).

Based on (1.11), we can explicitly construct a (normalized) hyperplane $H[\boldsymbol{u}_l, \tau_l]$ for each $l \in [N^-]$ that separates $\Phi(\boldsymbol{c}_l^-)$ and $\Phi(\mathcal{C}^+)$ with margin $\tilde{\mu} \asymp \delta^2 \lambda^{-1}$ (see Theorem 5.3). The resulting arrangement of the transformed sets $\Phi(\mathcal{C}^-)$ and $\Phi(\mathcal{C}^+)$ is illustrated in Figure 4. To conclude with the first layer, we need to ensure that the simplification achieved by Φ does not only apply to the center points but to the entire data set. Indeed, leveraging the *geometry-preserving* properties of random ReLU-layers (i.e., Φ preserves ℓ^2-distances between nearby points; see Theorem 4.1), it follows that $\Phi(\mathcal{X}_l^-)$ and $\Phi(\mathcal{X}^+)$ are also linearly separable for every $l \in [N^-]$ (still with margin $\tilde{\mu} \asymp \delta^2 \lambda^{-1}$, see Theorem 5.4). Hence, the geometric picture of Figure 4 remains true when replacing $\mathcal{C}^-$ and $\mathcal{C}^+$ by $\mathcal{X}^-$ and $\mathcal{X}^+$, respectively.[3] As for (1.11), this reasoning requires the bias $\boldsymbol{b}$ to be large enough (i.e., $\lambda \gtrsim \sqrt{\log(\lambda/\delta)}$) and exploits that it is *uniformly* distributed.

---

[3]Using such distance preservation properties of Φ is very different from a direct approach, according to which a random hyperplane $H[\boldsymbol{w}_i, b_i]$ would separate pairs of balls $(\mathcal{X}_l^-, \mathcal{X}_j^+)$. In fact, the latter event is much less likely than the separation of single points by a random hyperplane (this will become clear with Theorem 1.6 below). Instead, distance preservation is an integral property of the random layer, which exploits information from *all* coordinates of Φ.
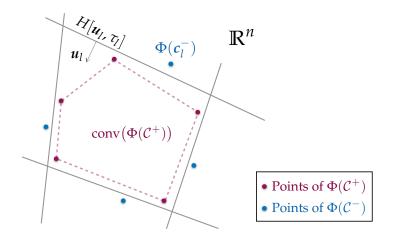
Figure 4: **The geometric effect of the first random ReLU-layer.** For each $l \in [N^-]$, the sets $\Phi(c_l^-)$ and $\Phi(\mathcal{C}^+)$ are linearly separable, or equivalently, it holds that $\Phi(\mathcal{C}^-) \cap \text{conv}(\Phi(\mathcal{C}^+)) = \varnothing$, where $\text{conv}(\cdot)$ is the convex hull operator. One can picture $\text{conv}(\Phi(\mathcal{C}^+))$ as a big "planet" which is orbited by small "satellites" namely the transformed center points in $\Phi(\mathcal{C}^-)$. For symmetry reasons, an analogous statement holds with high probability if the roles of $\mathcal{C}^-$ and $\mathcal{C}^+$ are interchanged.

Let us now turn to the second layer $\hat{\Phi} \colon \mathbb{R}^n \to \mathbb{R}^{\hat{n}}$, which builds directly on the geometric situation after applying Φ. It is again helpful to treat each coordinate individually:

$$[\hat{\Phi}(x)]_i = \sqrt{\tfrac{2}{\hat{n}}} \, \text{ReLU}(\langle \hat{w}_i, x \rangle + \hat{b}_i), \quad i = 1, \ldots, \hat{n}, \quad x \in \mathbb{R}^n,$$

where $\hat{w}_i \sim \mathsf{N}(0, I_n)$ and $\hat{b}_i$ is uniformly distributed on $[-\hat{\lambda}, \hat{\lambda}]$. Our main goal is to show that for every $l \in [N^-]$ there exist sufficiently many coordinates $i \in [\hat{n}]$ such that

$$[\hat{\Phi}(\Phi(\mathcal{X}_l^-))]_i \geq \tfrac{t}{\sqrt{\hat{n}}} \quad \text{and} \quad [\hat{\Phi}(\Phi(\mathcal{X}^+))]_i = 0, \tag{1.12}$$

where $t > 0$ depends on the complexity of $\mathcal{X}^-$ and $\mathcal{X}^+$; note that the vanishing coordinates are associated with data from $\mathcal{X}^+$ instead of $\mathcal{X}^-$. Our basic strategy to establish (1.12) is similar to (1.11), but there is a major difference: we now have to deal with the probability that a random hyperplane separates the sets $\hat{\mathcal{X}}_l^- := \Phi(\mathcal{X}_l^-)$ and $\hat{\mathcal{X}}^+ := \Phi(\mathcal{X}^+)$. The outcome of the first layer implies the existence of a separator, e.g., $H[u_l, \tau_l]$ (see also Figure 4), but this does not mean that it is likely to be found by a single random draw.[4] Certainly, the probability of successful separation may not only depend on the distance between $\hat{\mathcal{X}}_l^-$ and $\hat{\mathcal{X}}^+$, but also on their complexity and mutual arrangement. The following result makes this concern precise and forms a key component of our analysis; its proof can be found in Section 3.1. The notion of (ε, γ)-linear separability used below is formally introduced in Definition 3.1; for now, it is useful to think of a refinement of linear separability that captures how much a separating hyperplane for two sets can be perturbed, such that it is still separates the sets.

Theorem 1.6 *There exists an absolute constant $C > 0$ such that the following holds.*

For $\varepsilon \in [0,1]$, $\gamma > 0$, and $R \geq 1$, let $\mathcal{X}^-, \mathcal{X}^+ \subset R\mathbb{B}_2^d$ be two (ε, γ)-linearly separable sets[5] and put $\mu := (1 - \varepsilon)\gamma$. Let $g \sim \mathsf{N}(0, I_d)$ be a standard Gaussian vector and let τ be uniformly distributed on $[-\lambda, \lambda]$ for some $\lambda > 0$. For any $t \gtrsim w(\mathcal{X}^+ - \mathcal{X}^-) + R$ with $\lambda \gtrsim Rt\mu^{-1}$, the hyperplane $H[g, \tau]$ t-separates $\mathcal{X}^-$ from $\mathcal{X}^+$ with probability at least

$$\tfrac{t}{\lambda} \cdot \exp\left(-C \cdot t^2 \mu^{-2} \cdot \log(4(1 - \varepsilon)^{-1})\right). \tag{1.13}$$

---

[4]To be clear about this point, $H[u_l, \tau_l]$ does explicitly depend on the unknown sets $\mathcal{X}^-$ and $\mathcal{X}^+$. Hence, in contrast to random hyperplanes, it cannot be used for practical purposes offhand.

[5]Note that the notation for $\mathcal{X}^-$ and $\mathcal{X}^+$ is generic here. In the proof of our main result, we will apply Theorem 1.6 with $\mathcal{X}^- := \hat{\mathcal{X}}^+$ and $\mathcal{X}^+ := \hat{\mathcal{X}}_l^-$.

The central element of Theorem 1.6 is the mean width $w(\mathcal{X}^+ - \mathcal{X}^-)$, specifying the severity of the exponential decay in (1.13). An appropriate combination of Theorem 1.6 with a Chernoff bound will allow us to derive a statement of the form (1.12). With this at hand, it is then relatively straightforward to show that $F(\mathcal{X}^-) = \hat{\Phi}(\Phi(\mathcal{X}^-))$ and $F(\mathcal{X}^+) = \hat{\Phi}(\Phi(\mathcal{X}^+))$ are indeed linearly separable (cf. Corollary 5.2). Noteworthy is that the resulting margin and the number of required neurons $\hat{n}$ both inherit the exponential scaling from (1.13), which is reflected in all presented separation guarantees. This observation particularly explains why the mean width appears as a natural measure of complexity for the data sets.

To the best of our knowledge, Theorem 1.6 is a new result and could be of independent interest: it concerns the fundamental question of when pairs of sets are likely to be separated by a random hyperplane and when not. Perhaps not very surprisingly, the probability of success might scale poorly in the worst case, which is an inevitable consequence of the concentration of measure phenomenon. But for highly structured (low-dimensional) sets, the situation can be much more benign; finite point sets as considered in our analysis of the first layer are a good example.

1.5 Related Literature

The two-step separation procedure underlying our proofs (see Section 1.4) is inspired by a construction of An et al. [ABB15]. Their main result verifies that any two disjoint sets $\mathcal{X}^+, \mathcal{X}^- \subset \mathbb{R}^d$ can be made linearly separable by a *deterministic* two-layer NN. However, An et al. show a pure existence statement and their method is not feasible from an algorithmic perspective, since the selected weight parameters explicitly depend on the sets to be separated; furthermore, no informative bounds for the number of required neurons are provided. By using random weights and suitable notions of complexity (namely mutual covering), we are able to derive much more practical separation guarantees, which eliminate the aforementioned shortcomings. This achievement entails novel mathematical ingredients, most notably the separation capacity of random hyperplanes (Theorem 1.6) and uniform distance preservation by random ReLU layers (Theorem 4.1).

Below we will survey some works from the rich literature on random NNs that have notable conceptual similarities to our work. We are not aware of a comparable result that addresses the separation capacity of random NNs.

Approximation theory. A very active line of research investigates to what extent random neural networks are *universal approximators* (see e.g. [APVZ14; SGT18; YS19; NNSS20; HSSV21]) and the references therein). Specifically, one considers a class of real-valued functions on a domain in $\mathbb{R}^d$ (e.g., continuous or Lipschitz functions), an approximation metric (typically the L^2- or L^∞-norm), and a shallow neural network consisting of a ReLU layer with random weights followed by a linear layer with arbitrary weights (that may depend on the function to be approximated). The aforementioned works quantify which size of the random layer guarantees that the neural network can reach a pre-specified approximation error for every function in the given class. These results feature an exponential bottleneck, in the sense that the width of the random layer needs to scale exponentially in terms of the data dimension d to ensure accurate approximation (see also [NNSS20] for a refinement if the domain is a lower-dimensional smooth manifold).

In principle, one could try to approach Problem 1.1 by applying such an approximation result to a function $f : \mathbb{R}^d \to \mathbb{R}$ that takes values 1 and -1 on $\mathcal{X}^+$ and $\mathcal{X}^-$, respectively. However, existing approximation results cannot guarantee a *zero* approximation error on the two sets. In addition, even if an existing result would apply, it would lead to a more pessimistic result that would involve a random ReLU layer whose width scales exponentially in terms d, rather than the more refined, instance-specific complexity measures considered here.

Learning with random features. The random features method was introduced by Rahimi and Recht in [RR07] as a cheap computational alternative to kernel methods. The idea is to construct a random feature map so that inner products between random data features approximate kernel evaluations of the original data, provided that the feature dimension is high enough. A prime example are random Fourier features, which are designed to approximate the Gaussian kernel [RR07]. Instead of using a computationally expensive kernel method (e.g., kernel SVM), one can

instead use a linear method (e.g., SVM) on the random features. Several works have analyzed the generalization error of such random features methods and in particular quantified the feature dimension that guarantees a generalization performance that is on par with the associated kernel method ([RR08; Bac17; RR17; SGT18; LTOS21], see also [LHCS20] for a survey). Several of these results in particular apply when the random feature map is a random ReLU layer. Although these works indicate that random ReLU layers may transform data in a beneficial way for learning, they do not have a direct connection to Problem 1.1.

Random embeddings. A key component of our analysis is the capability of random ReLU-layers to preserve Euclidean distances with high probability (see Theorem 4.1). This finding is related to results on non-linear random embeddings, which play a major role in the field of quantized compressed sensing, see e.g. [JLBB13; PV14; OR15; CXJ17; DM18a; Dir19; DS20; XJ20]. In particular, our choice of the bias vector (see Definition 1.2) is inspired by *dithering*, a technique that has already proven useful in various signal reconstruction problems [JC17; DM18a; DM18b; XJ20; JMPS21]. A remarkable new (and somewhat counterintuitive) insight of the present work is that for appropriate non-linearities like the ReLU-activation, desirable distance preservation properties and data separation can be achieved simultaneously.

Theorem 4.1 is new in its own right and improves on a previous result by Giryes et al. [GSB16], see also [GSB20]. It is also closely related to work of Arpit and Bengio [AB19], who investigated the capability of a random ReLU layer as in Definition 1.2, but with bias $b = 0$, to preserve Euclidean norms.

Rare eclipse problem. Finally, we point out an interesting connection between the result on the separation capacity of random hyperplanes (Theorem 1.6) and the rare eclipse problem studied in [BMR17; CXJ17]. In both cases the goal is to map two linearly separable sets $\mathcal{X}^+, \mathcal{X}^- \subset \mathbb{R}^d$ into a lower dimensional space $\mathbb{R}^k$ using a random transform $T : \mathbb{R}^d \to \mathbb{R}^k$ such that $T(\mathcal{X}^+) \cap T(\mathcal{X}^-) = \varnothing$ with some probability p. In the rare eclipse problem, some arbitrarily small $\eta > 0$ is fixed in advance and the question is how small k can be so that $T(\mathcal{X}^+) \cap T(\mathcal{X}^-) = \varnothing$ holds with probability at least $p = 1 - \eta$. Using Gordon's escape through the mesh theorem [Gor88], the authors of [BMR17] showed that if $\mathcal{X}^+$ and $\mathcal{X}^-$ are disjoint closed and convex sets, then $k \gtrsim w^2(\text{cone}(\mathcal{X}^+ - \mathcal{X}^-) \cap \mathbb{S}^{d-1}) + \log(\eta^{-1})$ guarantees that $T(\mathcal{X}^+) \cap T(\mathcal{X}^-) = \varnothing$ holds with probability at least $1 - \eta$, where $T \in \mathbb{R}^{k \times d}$ is a standard Gaussian matrix. In comparison, Theorem 1.6 states that if $T(x) = \langle g, x \rangle + \tau$ for g a standard Gaussian random vector and $\tau \in [-\lambda, \lambda]$ uniformly distributed for $\lambda > 0$ large enough, then $T : \mathbb{R}^d \to \mathbb{R}$ maps two (ε, γ)-linearly separable sets $\mathcal{X}^+, \mathcal{X}^- \subset R\mathbb{B}_2^d$ to disjoint sets $T(\mathcal{X}^+), T(\mathcal{X}^+) \subset \mathbb{R}$ (which additionally have a certain minimal distance) with probability at least $R\lambda^{-1} \exp(-Cw^2(\mathcal{X}^+ - \mathcal{X}^-))$. Here, $C > 0$ is a constant that only depends on ε, γ and R. Hence, Theorem 1.6 achieves $T(\mathcal{X}^+) \cap T(\mathcal{X}^-) = \varnothing$ using a single coordinate ($k = 1$), however at the expense of a worse probability of success. Interestingly, both in the rare eclipse problem and Theorem 1.6, it is the squared Gaussian width of $\text{cone}(\mathcal{X}^+ - \mathcal{X}^-) \cap \mathbb{S}^{d-1}$, respectively of $\mathcal{X}^+ - \mathcal{X}^-$, which essentially determines the result.

1.6 Overview and Notation

The rest of the article is organized as follows: In Section 2, we present our main result, Theorem 2.3, based on the notion of mutual complexity (see Definition 2.1 and 2.2). The next two sections are then devoted to our main mathematical tools, namely separation by random hyperplanes (Section 3) and distance preservation (Section 4). Finally, the proof of Theorem 2.3 is given in Section 5, followed by a derivation of its variants (Theorem 1.3, 1.4, and 1.5) in Section 6.

Before proceeding, let us fix some standard notations and conventions that are commonly used in this paper. The letters c and C denote absolute (positive) constants, whose values may change from line to line. We speak of an *absolute constant* if its value does not depend on any other involved parameter. If an inequality holds up to an absolute constant C, we usually write $A \lesssim B$ instead of $A \leq C \cdot B$. The notation $A \asymp B$ is a shortcut for $A \lesssim B \lesssim A$.

For $d \in \mathbb{N}$, we set $[d] := \{1, \ldots, d\}$. The *cardinality* of an index set $I \subset [d]$ is denoted by $|I|$. Vectors and matrices are denoted by lower- and uppercase boldface letters, respectively. The *i*-th
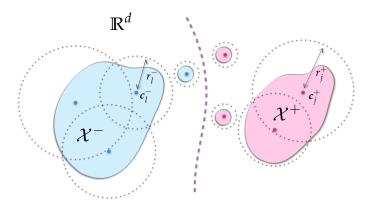
Figure 5: **Mutual covering.** This figure illustrates the geometric idea underlying Definition 2.1: those parts of $\mathcal{X}^-$ and $\mathcal{X}^+$ further away from the decision boundary may be covered by larger, and therefore fewer, Euclidean balls.

entry of a vector $z \in \mathbb{R}^d$ is denoted by $[z]_i$, or simply by z_i if there is no danger of confusion. We write $I_d \in \mathbb{R}^{d \times d}$ and $\mathbf{0} \in \mathbb{R}^d$ for the *identity matrix* and the *zero vector* in $\mathbb{R}^d$, respectively. For $1 \leq q \leq \infty$, we denote the ℓ^q-*norm* on $\mathbb{R}^d$ by $\| \cdot \|_q$ and the associated closed *unit ball* by $\mathbb{B}_q^d$. The *Euclidean unit sphere* is given by $\mathbb{S}^{d-1} := \{z \in \mathbb{R}^d \mid \|z\|_2 = 1\}$, and we also set $\mathbb{S}_+^{d-1} := \mathbb{S}^{d-1} \cap [0, \infty)^d$.

Let $\mathcal{X}, \mathcal{X}' \subset \mathbb{R}^d$ and $z \in \mathbb{R}^d$. The *linear cone* generated by $\mathcal{X}$ is denoted by $\operatorname{cone}(\mathcal{X}) := \{v\tilde{z} \mid \tilde{z} \in \mathcal{X}, v \geq 0\}$. The *Minkowski difference* between $\mathcal{X}$ and $\mathcal{X}'$ is defined by $\mathcal{X} - \mathcal{X}' := \{z_1 - z_2 \mid z_1 \in \mathcal{X}, z_2 \in \mathcal{X}'\}$, and we use the shortcut $\mathcal{X} - z := \mathcal{X} - \{z\}$. The *distance* between z and $\mathcal{X}$ is $\operatorname{dist}(z, \mathcal{X}) := \inf_{\tilde{z} \in \mathcal{X}} \|z - \tilde{z}\|_2$. Moreover, the *diameter* and *radius* of $\mathcal{X}$ are denoted by $\operatorname{diam}(\mathcal{X}) := \sup_{z_1, z_2 \in \mathcal{X}} \|z_1 - z_2\|_2$ and $\operatorname{rad}(\mathcal{X}) := \sup_{\tilde{z} \in \mathcal{X}} \|\tilde{z}\|_2$, respectively.

The L^q-*norm* of a real-valued random variable g is given by $\|g\|_{L^q} := (\mathbb{E}[|g|^q])^{1/q}$. We call g *sub-Gaussian* if $\|g\|_{\psi_2} := \inf \{v > 0 \mid \mathbb{E}[\exp(|g|^2/v^2)] \leq 2\} < \infty$; see [Ver18, Chap. 2 & 3] for more details on sub-Gaussian random variables and their properties. Finally, we write $g \sim \mathsf{N}(\mathbf{0}, I_d)$ if g is a *standard Gaussian random vector* in $\mathbb{R}^d$.

The *ceiling* and *floor function* of $z \in \mathbb{R}$ are denoted by $\lceil z \rceil$ and $\lfloor z \rfloor$, respectively.

2 Main Separation Result and Mutual Complexity

This section presents our most general solution to Problem 1.1, containing all guarantees from the introduction (Theorem 1.3, 1.4, and 1.5) as special cases. To formulate the main result, Theorem 2.3, we require two important definitions formalizing the idea of *mutual complexity* between two sets. The first one can be seen as a refinement of the uniform covering introduced in (1.8):

Definition 2.1 (Mutual covering) Let $\mathcal{X}^+, \mathcal{X}^- \subset \mathbb{R}^d$ and $\lambda > 0$. We call $\mathcal{C}^+ := \{c_1^+, \ldots, c_{N^+}^+\} \subset \mathbb{R}^d$ and $\mathcal{C}^- := \{c_1^-, \ldots, c_{N^-}^-\} \subset \mathbb{R}^d$ a λ-*mutual covering* for $\mathcal{X}^+$ and $\mathcal{X}^-$ if there exist $r_1^+, \ldots, r_{N^+}^+ \geq 0$ and $r_1^-, \ldots, r_{N^-}^- \geq 0$ such that

(i) the sets $\mathcal{X}_j^+ := \mathcal{X}^+ \cap \mathbb{B}_2^d(c_j^+, r_j^+)$ for $j \in [N^+]$, and $\mathcal{X}_l^- := \mathcal{X}^- \cap \mathbb{B}_2^d(c_l^-, r_l^-)$ for $l \in [N^-]$, cover $\mathcal{X}^+$ and $\mathcal{X}^-$, respectively;

(ii) $r_j^+ \leq \lambda^{-1} \operatorname{dist}^2(c_j^+, \mathcal{C}^-)$ for $j \in [N^+]$, and $r_l^- \leq \lambda^{-1} \operatorname{dist}^2(c_l^-, \mathcal{C}^+)$ for $l \in [N^-]$.

Furthermore, the sets $\mathcal{X}_1^+, \ldots, \mathcal{X}_{N^+}^+ \subset \mathcal{X}^+$ and $\mathcal{X}_1^-, \ldots, \mathcal{X}_{N^-}^- \subset \mathcal{X}^-$ are referred to as the *components* of the covering.

Although the notion of λ-mutual covering involves some technicalities, it is conceptually simple: We allow $\mathcal{X}^-$ and $\mathcal{X}^+$ to be covered by Euclidean balls of any radius, as long as the balls corresponding to different classes do not get too close in the sense of condition (ii). This constraint is

also consistent with the setting of Theorem 1.5, which is obtained by choosing $r_j^+ = r_l^- = c\delta^2/\lambda$. However, Definition 2.1 is much more flexible and accounts for the mutual arrangement of the classes. For example, those parts of $\mathcal{X}^-$ that are far away from the decision boundary may be covered by a few large balls, while smaller radii are only needed for data closer to $\mathcal{X}^+$; see Figure 5 for an illustration. In general, this strategy leads to more efficient coverings and motivates the following geometric complexity parameters:

Definition 2.2 (Mutual complexity) Let $\mathcal{X}^+, \mathcal{X}^- \subset \mathbb{R}^d$ and $\delta, \lambda > 0$. We say that $\mathcal{X}^+$ and $\mathcal{X}^-$ have (R, δ, λ)-*mutual complexity* (N^+, N^-, w^+, w^-) if there exists a λ-mutual covering $\mathcal{C}^+ = \{c_1^+, \ldots, c_{N^+}^+\}$ and $\mathcal{C}^- = \{c_1^-, \ldots, c_{N^-}^-\}$ for $\mathcal{X}^+$ and $\mathcal{X}^-$ such that

(i) $\max\limits_{j \in [N^+]} w(\mathcal{X}_j^+) \leq w^+$ and $\max\limits_{l \in [N^-]} w(\mathcal{X}_l^-) \leq w^-$;

(ii) $\mathcal{C}^+, \mathcal{C}^- \subset R\mathbb{B}_2^d$ are δ-separated.

It is useful to keep in mind that the covering numbers N^+ and N^- reflect the *global* size of $\mathcal{X}^+$ and $\mathcal{X}^-$, respectively, while w^+ and w^- should be viewed as *local* complexity measures. In contrast, the parameters (R, δ, λ) are not instance-specific and concern the general problem setting.

We are now ready to state the main result of this work:

Theorem 2.3 (Main result) *There exist absolute constants $C, C', C'' > 0$ such that the following holds.*
For $R \geq 1$, let $\mathcal{X}^-, \mathcal{X}^+ \subset R\mathbb{B}_2^d$ be δ-separated sets and let $\lambda \geq e\delta$ be such that $\lambda \gtrsim R\sqrt{\log(\lambda/\delta)}$. Furthermore, let $\mathcal{X}^+$ and $\mathcal{X}^-$ have $(R, \delta, C'\lambda)$-mutual complexity (N^+, N^-, w^+, w^-). We assume that $\Phi \colon \mathbb{R}^d \to \mathbb{R}^n$ and $\hat{\Phi} \colon \mathbb{R}^n \to \mathbb{R}^{\hat{n}}$ are two (independent) random ReLU-layers with maximal biases $\lambda, \hat{\lambda} \geq 0$, respectively, such that

$$n \gtrsim w^2(\mathrm{cone}(\mathcal{X}^- - \mathcal{X}^-) \cap \mathbb{S}^{d-1}) + w^2(\mathrm{cone}(\mathcal{X}^+ - \mathcal{X}^+) \cap \mathbb{S}^{d-1}),$$
$$n \gtrsim \left(\tfrac{\lambda}{\delta}\right)^8 \cdot \left(\lambda^{-2}(w^- + w^+)^2 + \log(2N^- N^+/\eta)\right) \tag{2.1}$$

and

$$\hat{\lambda} \gtrsim \left(\tfrac{\lambda}{\delta}\right)^4 \cdot (w^- + w(\mathcal{X}^+) + \lambda),$$
$$\hat{n} \gtrsim \left(\tfrac{\hat{\lambda}}{w^- + w(\mathcal{X}^+) + \lambda}\right) \cdot \exp\left(C \cdot (w^- + w(\mathcal{X}^+) + \lambda)^2 \cdot \lambda^6 \cdot \delta^{-8} \cdot \log(\lambda/\delta)\right) \cdot \log(N^-/\eta). \tag{2.2}$$

Then, given the two-layer random neural network $F \colon \mathbb{R}^d \to \mathbb{R}^{\hat{n}}$, $x \mapsto \hat{\Phi}(\Phi(x))$, with probability at least $1 - \eta$, the sets $F(\mathcal{X}^-), F(\mathcal{X}^+) \subset \hat{\lambda}\mathbb{B}_2^{\hat{n}}$ are linearly separable with margin

$$C'' \cdot \frac{(w^- + w(\mathcal{X}^+) + \lambda)^2}{\hat{\lambda}} \cdot \exp(-C(w^- + w(\mathcal{X}^+) + \lambda)^2 \cdot \lambda^6 \cdot \delta^{-8} \cdot \log(\lambda/\delta)). \tag{2.3}$$

Despite a strong resemblance to Theorem 1.5, the above result entails several important improvements. First, the exponential terms in (2.2) and (2.3) only depend on the localized mean width w^-, but not the covering number N^-. Hence, the global size of $\mathcal{X}^-$ does not have any (negative) impact here. The situation is different for $\mathcal{X}^+$, whose complexity is still captured by $w(\mathcal{X}^+)$. In fact, the following adaption of (1.10) clarifies the role of N^+:

$$w(\mathcal{X}^+) \lesssim w^+ + R\sqrt{\log N^+}. \tag{2.4}$$

The aforementioned asymmetry in Theorem 2.3 becomes especially useful when the set $\mathcal{X}^+$ is relatively "small" compared to $\mathcal{X}^-$. A prototypical example in this regard is a low-complexity set ($= \mathcal{X}^+$), say a small Euclidean ball, which is surrounded by a hypersphere ($= \mathcal{X}^-$); see Figure 6 for an illustration.

Another distinctive feature of Theorem 2.3 is the usage of *mutual* complexity. To understand its merits over the uniform covering considered in Theorem 1.5, it is worth revisiting the scenario of Figure 5: while the largest portion of the two classes is away from the (δ-separated) decision boundary, only a few "outliers" are close to it. Thus, a uniform covering would preset a very
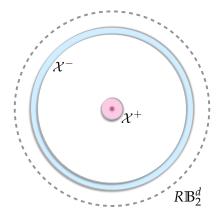
Figure 6: **An example of "asymmetric complexity" in the context of Theorem 2.3.** The set $\mathcal{X}^-$ corresponds to a thin hypersphere (say $\mathbb{S}^{d-1}$) around the origin, surrounding a small concentric ball $\mathcal{X}^+ = r\mathbb{B}_2^d$. If $r \lesssim 1/\sqrt{d}$, then w^- and $w(\mathcal{X}^+)$ are of constant order. Crucially, the covering number N^-, which scales exponentially in d, has no detrimental effect on the condition (2.2).

small radius (at the order $O(\lambda^{-1}\delta^2)$), which is appropriate for the outlier part but inefficient for the remaining bulk; this would lead to unnecessarily large covering numbers and thereby to poor complexity bounds (cf. (2.4)). In contrast, our mutual covering strategy is flexible enough to handle such data configurations. Therefore, Theorem 2.3 indeed presents an *instance-specific* solution to Problem 1.1, including a variant of outlier robustness.

Remark 2.4 (Possible extensions) For the sake of clarity, we have omitted some possible variations and generalizations of Theorem 2.3, which are however relatively straightforward to implement:

 (1) *Symmetry.* As discussed above, the asymmetric way of measuring complexity in Theorem 2.3 can be advantageous in certain situations. On the other hand, it is obvious that the roles of $\mathcal{X}^-$ and $\mathcal{X}^+$ are interchangeable. Hence, Theorem 2.3 could be "symmetrized" in this respect by a simple union bound argument.

 (2) *Non-linear activation.* The considered random network design is tailored to the ReLU-activation (see Definition 1.2). Nevertheless, our proof strategy is applicable to other functions as well, e.g., the thresholding activation. This might involve a slight adaption of Definition 2.1(ii) and lead to a different scaling of δ and λ in Theorem 2.3, but the qualitative statement remains valid.

$\Diamond$

3 Separation by Random Hyperplanes

The goal of this section is to prove Theorem 1.6, which is our main result on the separation of two sets by a random hyperplane. Before outlining the main steps of our proof, let us define the relevant notions of separability.

Definition 3.1 Let $\mathcal{X}^+, \mathcal{X}^- \subset \mathbb{R}^d$.

(a) Let $v \in \mathbb{R}^d \setminus \{0\}, \tau \in \mathbb{R}$ and $t \geq 0$. A hyperplane $H[v, \tau]$ *t-separates* $\mathcal{X}^-$ *from* $\mathcal{X}^+$ if

$$\langle v, x^- \rangle + \tau \leq -t \qquad \text{for all } x^- \in \mathcal{X}^-,$$
$$\langle v, x^+ \rangle + \tau > +t \qquad \text{for all } x^+ \in \mathcal{X}^+.$$

If $t = 0$, we simply say that $H[v, \tau]$ *separates* $\mathcal{X}^-$ *from* $\mathcal{X}^+$.

(b) Let $\varepsilon \in [0, 1]$ and $\gamma > 0$. We say that $\mathcal{X}^+$ and $\mathcal{X}^-$ are (ε, γ)-*linearly separable* if $\mathcal{X}^+$ and $\mathcal{X}^-$ are γ-separated (see Problem 1.1) and there exists $u \in \mathbb{S}^{d-1}$ such that

$$\langle u, x^+ - x^- \rangle \geq (1 - \varepsilon)\|x^+ - x^-\|_2 \qquad \text{for all } x^+ \in \mathcal{X}^+ \text{ and } x^- \in \mathcal{X}^-.$$

Recall from (1.1) that $\mathcal{X}^+$ and $\mathcal{X}^-$ are called *linearly separable with margin t* if they are t-separated by a hyperplane $H[v, \tau]$ with $\|v\|_2 = 1$. (ε, γ)-linearly separability is a strictly stronger condition (see also Proposition 3.2(iii) below). Intuitively, it captures how much a separating hyperplane can be perturbed, such that it still separates the sets $\mathcal{X}^+$ and $\mathcal{X}^-$; geometrically, the parameter ε controls the narrowness of $\mathrm{cone}(\mathcal{X}^+ - \mathcal{X}^-)$.

Proof sketch for Theorem 1.6: By a rescaling argument, we can assume that $R = 1$. For a $k \in \mathbb{N}$ specified below, we represent the standard Gaussian vector $g \in \mathbb{R}^d$ by $g = G^T V'$, where $G \in \mathbb{R}^{k \times d}$ is a standard Gaussian matrix, $V' \in \mathbb{S}^{k-1}$ is uniformly distributed and G, V' are independent. We then observe that, for any $\rho \geq 0$, the hyperplane $H[g, \tau]$ ρ-separates $\mathcal{X}^-$ from $\mathcal{X}^+$ if and only if the hyperplane $H[\sqrt{k} V', \tau]$ ρ-separates $\frac{1}{\sqrt{k}} G \mathcal{X}^-$ from $\frac{1}{\sqrt{k}} G \mathcal{X}^+$. Therefore, one can prove Theorem 1.6 by first showing that for k large enough, the sets $\frac{1}{\sqrt{k}} G \mathcal{X}^-$ and $\frac{1}{\sqrt{k}} G \mathcal{X}^+$ are again linearly separable with constant probability and second, showing that conditioned on this event the hyperplane $H[\sqrt{k} V', \tau]$ ρ-separates $\frac{1}{\sqrt{k}} G \mathcal{X}^-$ and $\frac{1}{\sqrt{k}} G \mathcal{X}^+$ with probability p, where ρ and p are specified in Theorem 1.6. Specifically, the main technical steps are

1. to show that if
$$k \gtrsim \gamma^{-2} (1 - \varepsilon)^{-2} (w^2(\mathcal{X}^+ - \mathcal{X}^-) + 1),$$
then the linear transformation $\frac{1}{\sqrt{k}} G$ maps the (ε, γ)-linearly separable sets $\mathcal{X}^-$ and $\mathcal{X}^+$ to $(\frac{1+\varepsilon}{2}, \frac{\gamma}{2})$-linearly separable sets $\frac{1}{\sqrt{k}} G \mathcal{X}^-$ and $\frac{1}{\sqrt{k}} G \mathcal{X}^+$ with probability at least $\frac{1}{2}$.

2. to derive a general separation result for two (ε, γ)-linearly separable sets by a random hyperplane $H[V, \tau]$, where V is uniformly distributed on Euclidean sphere (Theorem 3.3 and Corollary 3.4).

Let us now give the proof in full detail. We start with a simple proposition that relates our notions of separability.

Proposition 3.2 *Let $\mathcal{X}^+, \mathcal{X}^- \subset \mathbb{R}^d$. The following relationships hold:*

(i) *If $\mathcal{X}^+$ and $\mathcal{X}^-$ are linearly separable with margin μ, then they are 2μ-separated.*

(ii) *If a hyperplane $H[u, \tau]$ t-separates $\mathcal{X}^-$ from $\mathcal{X}^+$, then $\mathcal{X}^+$ and $\mathcal{X}^-$ are linearly separable with margin $t / \|u\|_2$.*

(iii) *If $\mathcal{X}^+$ and $\mathcal{X}^-$ are (ε, γ)-linearly separable, then they are linearly separable with margin $\frac{(1-\varepsilon)\gamma}{2}$.*

(iv) *If $\mathcal{X}^+$ and $\mathcal{X}^-$ are linearly separable with margin μ and $\mathrm{diam}(\mathcal{X}^+ - \mathcal{X}^-) \leq R$, then they are $(\frac{R-2\mu}{R}, 2\mu)$-linearly separable.*

Proof. To show (i), observe that by assumption there exist $u \in \mathbb{S}^{d-1}$ and $\tau \in \mathbb{R}$ such that

$$\langle u, x^- \rangle + \tau \leq -\mu \qquad \text{for all } x^- \in \mathcal{X}^-,$$
$$\langle u, x^+ \rangle + \tau \geq +\mu \qquad \text{for all } x^+ \in \mathcal{X}^+.$$

It follows that $\langle u, x^+ - x^- \rangle \geq 2\mu$ for all $x^+ \in \mathcal{X}^+, x^- \in \mathcal{X}^-$. By the Cauchy-Schwarz inequality, $\langle u, x^+ - x^- \rangle \leq \|x^+ - x^-\|_2$, which shows the claim. For (ii), it suffices to note that if $H[u, \tau]$ t-separates $\mathcal{X}^-$ from $\mathcal{X}^+$, then $H[\frac{u}{\|u\|_2}, \frac{\tau}{\|u\|_2}]$ $\frac{t}{\|u\|_2}$-separates $\mathcal{X}^-$ from $\mathcal{X}^+$. Let us next show (iii). If $\mathcal{X}^+$ and $\mathcal{X}^-$ are (ε, γ)-linearly separable, then there exists $u \in \mathbb{S}^{d-1}$ such that

$$\langle u, x^+ - x^- \rangle \geq (1 - \varepsilon)\gamma \qquad \text{for all } x^+ \in \mathcal{X}^+ \text{ and } x^- \in \mathcal{X}^-.$$

Set $\tau = -\frac{(1-\varepsilon)\gamma}{2} + \inf_{x^- \in \mathcal{X}^-} \langle u, -x^- \rangle$. Fix $x' \in \mathcal{X}^+$. Since

$$\inf_{x^- \in \mathcal{X}^-} \langle u, -x^- \rangle = -\langle u, x' \rangle + \inf_{x^- \in \mathcal{X}^-} \langle u, x' - x^- \rangle \geq -\langle u, x' \rangle + (1 - \varepsilon)\gamma,$$

we see that $\tau \in \mathbb{R}$. Further, for any $x^- \in \mathcal{X}^-$,

$$\langle u, x^- \rangle + \tau = \langle u, x^- \rangle - \frac{(1-\varepsilon)\gamma}{2} + \inf_{x^- \in \mathcal{X}^-} \langle u, -x^- \rangle$$

$$= \langle u, x^- \rangle - \sup_{x^- \in \mathcal{X}^-} \langle u, x^- \rangle - \tfrac{(1-\varepsilon)\gamma}{2} \leq -\tfrac{(1-\varepsilon)\gamma}{2}$$

and for any $x^+ \in \mathcal{X}^+$,

$$\langle u, x^+ \rangle + \tau = \langle u, x^+ \rangle - \tfrac{(1-\varepsilon)\gamma}{2} + \inf_{x^- \in \mathcal{X}^-} \langle u, -x^- \rangle$$
$$= \inf_{x^- \in \mathcal{X}^-} \langle u, x^+ - x^- \rangle - \tfrac{(1-\varepsilon)\gamma}{2} \geq \tfrac{(1-\varepsilon)\gamma}{2}.$$

Since $\|u\|_2 = 1$, it follows that the hyperplane $H[u, \tau]$ linearly separates $\mathcal{X}^-$ and $\mathcal{X}^+$ with margin $\tfrac{(1-\varepsilon)\gamma}{2}$. Finally, let us show (iv). By (i) we know that $\mathcal{X}^+$ and $\mathcal{X}^-$ are 2μ-separated. Let $u \in \mathrm{S}^{d-1}$ and $\tau \in \mathbb{R}$ be such that

$$\langle u, x^- \rangle + \tau \leq -\mu \qquad \text{for all } x^- \in \mathcal{X}^-,$$
$$\langle u, x^+ \rangle + \tau \geq +\mu \qquad \text{for all } x^+ \in \mathcal{X}^+.$$

It follows that $\langle u, x^+ - x^- \rangle \geq 2\mu$ for all $x^+ \in \mathcal{X}^+, x^- \in \mathcal{X}^-$. Since $\mathrm{diam}(\mathcal{X}^+ - \mathcal{X}^-) \leq R$, we also have $2\mu \geq \tfrac{2\mu}{R} \|x^+ - x^-\|_2$ for all $x^+ \in \mathcal{X}^+, x^- \in \mathcal{X}^-$. Together this yields

$$\langle u, x^+ - x^- \rangle \geq \left(1 - \tfrac{R - 2\mu}{R}\right) \|x^+ - x^-\|_2 \qquad \text{for all } x^+ \in \mathcal{X}^+ \text{ and } x^- \in \mathcal{X}^-.$$

∎

The next result gives a lower bound for the probability that a random hyperplane $H[V, \tau]$ separates two (ε, γ)-linearly separable sets $\mathcal{X}^+, \mathcal{X}^- \subset R\mathbb{B}_2^d$, if $\tau \in [-\lambda, \lambda]$ is uniformly distributed for $\lambda \geq R$ and V is uniformly distributed on S^{d-1}. As detailed in our proof sketch, this result (more precisely, Corollary 3.4) forms a crucial ingredient of our proof of Theorem 1.6.

Theorem 3.3 *There exist absolute constants $c, C > 0$ such that the following holds. For $\varepsilon \in [0, 1]$ and $\gamma > 0$, consider (ε, γ)-linearly separable sets $\mathcal{X}^+, \mathcal{X}^- \subset R\mathbb{B}_2^d$. Let $V \in \mathrm{S}^{d-1}$ and $\tau \in [-\lambda, \lambda]$ be both uniformly distributed. If $\lambda \geq R$, then with probability at least*

$$c \tfrac{\gamma}{\lambda} (1 - \varepsilon)(\sqrt{\varepsilon} + \tfrac{1}{\sqrt{d}}) \exp(-C\varepsilon d \log(2(1-\varepsilon)^{-1})),$$

the hyperplane $H[V, \tau]$ t-separates $\mathcal{X}^-$ from $\mathcal{X}^+$ with $t = c\gamma(1 - \varepsilon)(\sqrt{\varepsilon} + \tfrac{1}{\sqrt{d}})$.

Corollary 3.4 *There exist absolute constants $c, C > 0$ such that the following holds. For $\varepsilon \in [0, 1]$ and $\gamma > 0$, consider (ε, γ)-linearly separable sets $\mathcal{X}^+, \mathcal{X}^- \subset R\mathbb{B}_2^d$. Let $\nu > 0$. Let $V \in \mathrm{S}^{d-1}$ and $\tau \in [-\lambda, \lambda]$ be both uniformly distributed. If $\lambda \geq \nu R$, then with probability at least*

$$c \tfrac{\nu\gamma}{\lambda} (1 - \varepsilon)(\sqrt{\varepsilon} + \tfrac{1}{\sqrt{d}}) \exp(-C\varepsilon d \log(2(1-\varepsilon)^{-1})),$$

the hyperplane $H[\nu V, \tau]$ t-separates $\mathcal{X}^-$ from $\mathcal{X}^+$ with $t = c\gamma(1 - \varepsilon)(\sqrt{\varepsilon} + \tfrac{1}{\sqrt{d}})\nu$.

Proof. For $t \geq 0$ the hyperplane $H[\nu V, \tau]$ $t\nu$-separates $\mathcal{X}^-$ from $\mathcal{X}^+$ if and only if the hyperplane $H[V, \tfrac{1}{\nu}\tau]$ t-separates $\mathcal{X}^-$ from $\mathcal{X}^+$. The random variable $\tau' := \tfrac{1}{\nu}\tau$ is uniformly distributed on $[-\lambda', \lambda']$ for $\lambda' = \tfrac{\lambda}{\nu}$. The result follows from Theorem 3.3 for $t = c\gamma(1 - \varepsilon)(\sqrt{\varepsilon} + \tfrac{1}{\sqrt{d}})$. ∎

For the proof of Theorem 3.3 we need the following standard result (see, e.g., [BLM13, Section 7.2]), which precisely describes the surface measure of a spherical cap.

Lemma 3.5 *Let $V \in \mathrm{S}^{d-1}$ be uniformly distributed. Let $\delta \in (0, 1]$ and $d \geq 2\delta^{-2}$. For any $u \in \mathrm{S}^{d-1}$,*

$$\tfrac{1}{6\delta\sqrt{d}} (1 - \delta^2)^{\frac{d-1}{2}} \leq \mathbb{P}\left(\langle V, u \rangle \geq \delta\right) \leq \tfrac{1}{2\delta\sqrt{d}} (1 - \delta^2)^{\frac{d-1}{2}}.$$

If additionally $\delta \leq \tfrac{1}{\sqrt{2}}$, then

$$\mathbb{P}\left(\langle V, u \rangle \geq \delta\right) \geq \tfrac{1}{2} \exp(-2\delta^2 d).$$

Proof of Theorem 3.3. For $\theta \geq 0$ define the event

$$A_V(\theta) := \left\{ \inf_{x^+ \in \mathcal{X}^+, x^- \in \mathcal{X}^-} \langle V, x^+ - x^- \rangle \geq \theta \right\}.$$

For any $s \geq 0$,

$$\mathbb{P}(H[V, \tau] \ s\text{-separates } \mathcal{X}^- \text{ from } \mathcal{X}^+)$$
$$\geq \mathbb{P}(\{H[V, \tau] \ s\text{-separates } \mathcal{X}^- \text{ from } \mathcal{X}^+\} \cap A_V(\theta))$$
$$= p_1(s, \theta) \cdot p_2(\theta),$$

where

$$p_1(s, \theta) := \mathbb{P}(H[V, \tau] \ s\text{-separates } \mathcal{X}^- \text{ from } \mathcal{X}^+ \mid A_V(\theta)) \quad \text{and} \quad p_2(\theta) := \mathbb{P}(A_V(\theta)).$$

Next, we bound both factors $p_1(s, \theta)$ and $p_2(\theta)$ from below.
Lower bound for $p_1(s, \theta)$: Let us show that if $\lambda \geq R$, then for $s \leq \frac{\theta}{2}$,

$$\mathbb{P}(H[V, \tau] \ s\text{-separates } \mathcal{X}^- \text{ from } \mathcal{X}^+ \mid A_V(\theta)) \geq \frac{\theta - 2s}{2\lambda}.$$

If $\lambda \geq R$, then $a := \sup_{x^- \in \mathcal{X}^-} \langle V, x^- \rangle \in [-\lambda, \lambda]$ and $b := \inf_{x^+ \in \mathcal{X}^+} \langle V, x^+ \rangle \in [-\lambda, \lambda]$. Further, on the event $A_V(\theta)$ it holds $b - a \geq \theta$. Let $s \leq \frac{\theta}{2}$. If $-\tau \in [a + s, b - s]$, then

$$\langle V, x^- \rangle + \tau \leq -s \qquad \text{for all } x^- \in \mathcal{X}^-$$
$$\langle V, x^+ \rangle + \tau \geq +s \qquad \text{for all } x^+ \in \mathcal{X}^+.$$

Since $\mathbb{P}_\tau(-\tau \in [a + s, b - s]) \geq \frac{\theta - 2s}{2\lambda}$, this shows

$$\mathbb{P}_\tau(H[V, \tau] \ s\text{-separates } \mathcal{X}^- \text{ from } \mathcal{X}^+ \mid A_V(\theta)) \geq \frac{\theta - 2s}{2\lambda}.$$

For $s = \frac{\theta}{4}$ we obtain

$$\mathbb{P}(H[V, \tau] \ \tfrac{\theta}{4}\text{-separates } \mathcal{X}^- \text{ from } \mathcal{X}^+) \geq \frac{\theta}{4\lambda} \cdot p_2(\theta). \tag{3.1}$$

Lower bound for $p_2(\theta)$: Set $\mathcal{X} := \text{cone}(\mathcal{X}^+ - \mathcal{X}^-) \cap \mathbb{S}^{d-1}$ and for $\delta \in [0, 1]$ define

$$\alpha_\varepsilon(\delta) := \max\{\delta - \sqrt{2\varepsilon}, 1 - \varepsilon - \sqrt{2}\sqrt{1 - \delta}\}.$$

Since $\mathcal{X}^+$ and $\mathcal{X}^-$ are (ε, γ)-linearly separable there exists $u \in \mathbb{S}^{d-1}$ such that

$$\langle u, z \rangle \geq 1 - \varepsilon \qquad \text{for all } z \in \mathcal{X}. \tag{3.2}$$

Let us show that for any $\delta \in [0, 1]$ with $\alpha_\varepsilon(\delta) \geq 0$,

$$\{\langle V, u \rangle \geq \delta\} \subset A_V(\alpha_\varepsilon(\delta)\gamma). \tag{3.3}$$

First observe that for any $\delta \in [0, 1]$,

$$\{\langle V, u \rangle \geq \delta\} \subset \left\{ \inf_{z \in \mathcal{X}} \langle V, z \rangle \geq \alpha_\varepsilon(\delta) \right\}. \tag{3.4}$$

Indeed, by (3.2), $\|z - u\|_2 \leq \sqrt{2\varepsilon}$ for every $z \in \mathcal{X}$. Since $V \in \mathbb{S}^{d-1}$ it follows

$$\langle V, z \rangle \geq \langle V, u \rangle - \|z - u\|_2 \geq \langle V, u \rangle - \sqrt{2\varepsilon} \geq \delta - \sqrt{2\varepsilon} \quad \text{for all } z \in \mathcal{X}, \tag{3.5}$$

if $\langle V, u \rangle \geq \delta$. Moreover, $\langle V, u \rangle \geq \delta \Leftrightarrow \|V - u\|_2 \leq \sqrt{2}\sqrt{1 - \delta}$. Therefore, if $\langle V, u \rangle \geq \delta$ then for every $z \in \mathcal{X}$,

$$\langle V, z \rangle \geq \langle u, z \rangle - \|V - u\|_2 \geq 1 - \varepsilon - \sqrt{2}\sqrt{1 - \delta}. \tag{3.6}$$

Inequalities (3.5) and (3.6) imply (3.4). If $\inf_{z \in \mathcal{X}} \langle V, z \rangle \geq \alpha_\varepsilon(\delta)$ and $\alpha_\varepsilon(\delta) \geq 0$, the following holds for every $x^+ \in \mathcal{X}^+, x^- \in \mathcal{X}^-$:

$$\langle V, x^+ - x^- \rangle \geq \alpha_\varepsilon(\delta) \|x^+ - x^-\|_2 \geq \alpha_\varepsilon(\delta)\gamma,$$

where for the second inequality we used that $\alpha_\varepsilon(\delta) \geq 0$ and that $\mathcal{X}^+$ and $\mathcal{X}^-$ are γ-separated. In combination with (3.4) this shows (3.3). From (3.3) it follows that for every $\delta \in [0, 1]$ with $\alpha_\varepsilon(\delta) \geq 0$ and every $\theta \leq \alpha_\varepsilon(\delta)\gamma$,

$$p_2(\theta) \geq \mathbb{P}(\langle V, u \rangle \geq \delta). \tag{3.7}$$

In order to bound the probability on the right hand side from below, we distinguish the cases $\varepsilon \leq \frac{1}{32}$ and $\varepsilon > \frac{1}{32}$.

Case $\varepsilon \leq \frac{1}{32}$: If $\delta \in [0, 1]$ satisfies $\delta \geq \sqrt{8\varepsilon}$, then $\delta - \sqrt{2\varepsilon} \geq \frac{\delta}{2}$, which implies $\alpha_\varepsilon(\delta) \geq \frac{\delta}{2} \geq 0$. Therefore, by (3.7) the following holds: For all $\delta \in [\sqrt{8\varepsilon}, 1]$,

$$p_2(\tfrac{\delta}{2}\gamma) \geq \mathbb{P}(\langle V, u \rangle \geq \delta).$$

By Lemma 3.5 we obtain that for all $\delta \in (\sqrt{8\varepsilon}, \frac{1}{\sqrt{2}}]$ with $d \geq 2\delta^{-2}$,

$$p_2(\tfrac{\delta}{2}\gamma) \geq \tfrac{1}{2}\exp(-2\delta^2 d).$$

Applying (3.1) for $\theta = \frac{\delta}{2}\gamma$, we obtain that for all $\delta \in (\sqrt{8\varepsilon}, \frac{1}{\sqrt{2}}]$ with $d \geq 2\delta^{-2}$,

$$\mathbb{P}(H[V, \tau] \tfrac{\delta\gamma}{8}\text{-separates } \mathcal{X}^- \text{ from } \mathcal{X}^+) \geq \tfrac{\delta\gamma}{8\lambda} \cdot \tfrac{1}{2}\exp(-2\delta^2 d).$$

The choice $\delta = \sqrt{8\varepsilon} + \sqrt{\frac{2}{d}}$, which satisfies $\delta \in (\sqrt{8\varepsilon}, \frac{1}{\sqrt{2}}]$ and $d \geq 2\delta^{-2}$, yields

$$\mathbb{P}(H[V, \tau] \tfrac{\gamma}{8}(\sqrt{8\varepsilon} + \sqrt{\tfrac{2}{d}})\text{-separates } \mathcal{X}^- \text{ from } \mathcal{X}^+) \geq c\tfrac{\gamma}{\lambda}(\sqrt{\varepsilon} + \tfrac{1}{\sqrt{d}}) \cdot \exp(-C\varepsilon d)$$

for absolute constants $c, C > 0$. The result in the case $\varepsilon \leq \frac{1}{32}$ follows by observing that $1 - \varepsilon \sim 1$.

Case $\varepsilon > \frac{1}{32}$: Set $\delta' = 1 - \frac{1}{4}(1 - \varepsilon)^2$. Then $\delta' \in [0, 1]$ and $\alpha_\varepsilon(\delta') \geq 1 - \varepsilon - \sqrt{2}\sqrt{1 - \delta'} = (1 - \frac{1}{\sqrt{2}})(1 - \varepsilon) \geq \frac{1}{4}(1 - \varepsilon) \geq 0$. Therefore, by (3.7) in combination with Lemma 3.5 we obtain

$$p_2(\tfrac{1}{4}(1 - \varepsilon)\gamma) \geq \tfrac{1}{6\delta'\sqrt{d}}(1 - \delta'^2)^{\frac{d-1}{2}}.$$

Observe that

$$\begin{aligned}
\tfrac{1}{6\delta'\sqrt{d}}(1 - \delta'^2)^{\frac{d-1}{2}} &\geq \tfrac{1}{6\sqrt{d}}\exp(-d\log((1 - \delta'^2)^{-1})) \\
&\geq \tfrac{1}{6\sqrt{d}}\exp(-d\log((1 - \delta')^{-1})) \\
&= \tfrac{1}{6\sqrt{d}}\exp(-2d\log(2(1 - \varepsilon)^{-1})) \\
&\geq \exp(-3d\log(2(1 - \varepsilon)^{-1})).
\end{aligned}$$

Applying (3.1) for $\theta = \frac{1}{4}(1 - \varepsilon)\gamma$, we obtain

$$\mathbb{P}(H[V, \tau] \tfrac{(1-\varepsilon)\gamma}{16}\text{-separates } \mathcal{X}^- \text{ from } \mathcal{X}^+) \geq \tfrac{(1-\varepsilon)\gamma}{16\lambda} \cdot \exp(-3d\log(2(1 - \varepsilon)^{-1})).$$

The result in the case $\varepsilon > \frac{1}{32}$ follows by observing that $\sqrt{\varepsilon} + \frac{1}{\sqrt{d}} \sim 1$ and $\varepsilon \sim 1$. This completes the proof. ∎

3.1 Proof of Theorem 1.6

For the proof of Theorem 1.6 we need one final ingredient, namely Lemma 3.6. It gives a sufficient condition under which a set, which is contained in a spherical cone, is again contained in a spherical cone after a linear transformation. Using this lemma, we will show that a Gaussian matrix with enough rows maps (ε, γ)-linear separable sets to $(\frac{1+\varepsilon}{2}, \frac{\gamma}{2})$-linearly separable sets with constant probability (step 1 of the proof sketch at the beginning of Section 3).

Lemma 3.6 *Let $u \in \mathbb{S}^{d-1}$, $t \in [0, 1]$ and $\mathcal{X} \subset R\mathbb{B}_2^d$ satisfy*

$$\langle u, x \rangle \geq t\|x\|_2 \qquad \text{for all } x \in \mathcal{X}.$$

For $\kappa \in (0, \frac{1}{2}]$, $\alpha, \beta \in [0, R]$ and $A \in \mathbb{R}^{k \times d}$ assume that the following hold:

1. $(1 - \kappa)\|u\|_2^2 \leq \|Au\|_2^2 \leq (1 + \kappa)\|u\|_2^2$,

2. *For all $x \in \mathcal{X}$,*
$$\|u - x\|_2 - \alpha \leq \|Au - Ax\|_2 \leq \|u - x\|_2 + \alpha,$$

3. *For all $x \in \mathcal{X}$,*
$$\|x\|_2 - \beta \leq \|Ax\|_2 \leq \|x\|_2 + \beta.$$

Then $A\mathcal{X} \subset 2R\mathbb{B}_2^k$ and

$$\langle \tfrac{Au}{\|Au\|_2}, Ax \rangle \geq \tfrac{t}{\sqrt{1+\kappa}}\|Ax\|_2 - \tfrac{\kappa}{\sqrt{2}} - \sqrt{2}(\tfrac{3}{2}R + t)\beta - \tfrac{3}{\sqrt{2}}(1 + R)\alpha \qquad \text{for all } x \in \mathcal{X}. \tag{3.8}$$

Proof. For any $x \in \mathcal{X}$,

$$\left|\|Ax\|_2^2 - \|x\|_2^2\right| = \left|\|Ax\|_2 - \|x\|_2\right| \cdot \left|\|Ax\|_2 + \|x\|_2\right| \leq \beta(2\|x\|_2 + \beta) \leq 3R\beta.$$

Analogously,

$$\left|\|Au - Ax\|_2^2 - \|u - x\|_2^2\right| \leq 3(1 + R)\alpha.$$

Therefore,

$$\begin{aligned}
\langle Au, Ax \rangle &= \tfrac{1}{2}\|Au\|_2^2 + \tfrac{1}{2}\|Ax\|_2^2 - \tfrac{1}{2}\|A(u - x)\|_2^2 \\
&\geq \tfrac{1}{2}(1 - \kappa)\|u\|_2^2 + \tfrac{1}{2}\|x\|_2^2 - \tfrac{3}{2}R\beta - \tfrac{1}{2}\|u - x\|_2^2 - \tfrac{3}{2}(1 + R)\alpha \\
&= \langle u, x \rangle - \tfrac{\kappa}{2}\|u\|_2^2 - \tfrac{3}{2}R\beta - \tfrac{3}{2}(1 + R)\alpha \\
&\geq t\|x\|_2 - \tfrac{\kappa}{2} - \tfrac{3}{2}R\beta - \tfrac{3}{2}(1 + R)\alpha \\
&\geq t\|Ax\|_2 - \tfrac{\kappa}{2} - (\tfrac{3}{2}R + t)\beta - \tfrac{3}{2}(1 + R)\alpha.
\end{aligned}$$

Using $\tfrac{1}{\sqrt{2}} \leq \|Au\|_2 \leq \sqrt{1 + \kappa}$, we obtain (3.8). ∎

We are ready to prove our main result on the separation of two sets by a random hyperplane:

Proof of Theorem 1.6. First note that it suffices to prove the result for $R = 1$. Indeed, the general result then follows by a rescaling argument.

Let $k \in \mathbb{N}$, let $V' \in \mathbb{S}^{k-1}$ be uniformly distributed, $G \in \mathbb{R}^{k \times d}$ a standard Gaussian matrix and $\tau \in [-\lambda, \lambda]$ be uniformly distributed. Let all random variables be independent. We calculate the characteristic function of $g := G^T V' \in \mathbb{R}^d$ to check that g is standard Gaussian. For any $z \in \mathbb{R}^d$,

$$\varphi_g(z) = \mathbb{E}\big[e^{i\langle z, g \rangle}\big] = \mathbb{E}\big[e^{i\langle Gz, V' \rangle}\big] = \mathbb{E}\big[e^{i\langle G\frac{z}{\|z\|_2}, \|z\|_2 V' \rangle}\big].$$

Since $G\frac{z}{\|z\|_2}$ is standard Gaussian, $\mathbb{E}\big[e^{i\langle G\frac{z}{\|z\|_2}, x \rangle}\big] = e^{-\frac{\|x\|_2^2}{2}}$ for any $x \in \mathbb{R}^d$. Therefore, using the independence of G and V', we obtain

$$\varphi_g(z) = \mathbb{E}_{V'}\mathbb{E}_G\big[e^{i\langle G\frac{z}{\|z\|_2}, \|z\|_2 V' \rangle}\big] = \mathbb{E}_{V'}\big[e^{-\frac{\|\|z\|_2 V'\|_2^2}{2}}\big] = e^{-\frac{\|z\|_2^2}{2}},$$

where we used that $\|V'\|_2 = 1$. This shows that g is standard Gaussian. Set $A = \frac{1}{\sqrt{k}}G$. For $\varepsilon' \in [0,1]$ and $\gamma' \geq 0$ define the event

$$\mathsf{A}_{\varepsilon',\gamma'} := \{A\mathcal{X}^+, A\mathcal{X}^- \subset 5\mathbb{B}_2^k \text{ are } (\varepsilon', \gamma')\text{-linearly separable}\}.$$

Using that $g = A^T\sqrt{k}V'$, we obtain for any $\rho \geq 0$,

$$\begin{aligned}
&\mathbb{P}(H[g,\tau] \ \rho\text{-separates } \mathcal{X}^- \text{ from } \mathcal{X}^+) && (3.9)\\
&\mathbb{P}(H[\sqrt{k}V',\tau] \ \rho\text{-separates } A\mathcal{X}^- \text{ from } A\mathcal{X}^+)\\
&\geq \mathbb{P}(\{H[\sqrt{k}V',\tau] \ \rho\text{-separates } A\mathcal{X}^- \text{ from } A\mathcal{X}^+\} \cap \mathsf{A}_{\varepsilon',\gamma'})\\
&= \mathbb{P}(H[\sqrt{k}V',\tau] \ \rho\text{-separates } A\mathcal{X}^- \text{ from } A\mathcal{X}^+ \mid \mathsf{A}_{\varepsilon',\gamma'}) \cdot \mathbb{P}(\mathsf{A}_{\varepsilon',\gamma'}).
\end{aligned}$$

Since $\mathcal{X}^+$ and $\mathcal{X}^-$ are (ε,γ)-linearly separable, there exists a vector $u \in \mathbb{S}^{d-1}$ such that

$$\langle u, x^+ - x^- \rangle \geq (1-\varepsilon)\|x^+ - x^-\|_2 \qquad \text{for all } x^+ \in \mathcal{X}^+, x^- \in \mathcal{X}^-.$$

For $\kappa \in (0, \frac{1}{2}], \alpha \in [0,2], \beta \in [0, \frac{\gamma}{2}]$ define the event $\mathsf{B}_{\kappa,\alpha,\beta}$ where:

1. $(1-\kappa)\|u\|_2^2 \leq \|Au\|_2^2 \leq (1+\kappa)\|u\|_2^2$,

2. For all $x^+ \in \mathcal{X}^+, x^- \in \mathcal{X}^-$,

$$\|u - (x^+ - x^-)\|_2 - \alpha \leq \|Au - A(x^+ - x^-)\|_2 \leq \|u - (x^+ - x^-)\|_2 + \alpha,$$

3. For all $x^+ \in \mathcal{X}^+, x^- \in \mathcal{X}^-$,

$$\|x^+ - x^-\|_2 - \beta \leq \|A(x^+ - x^-)\|_2 \leq \|x^+ - x^-\|_2 + \beta,$$

4. There exists $x^- \in \mathcal{X}^-$ such that $\|Ax^-\|_2^2 \leq (1+\kappa)\|x^-\|_2^2$,

5. There exists $x^+ \in \mathcal{X}^+$ such that $\|Ax^+\|_2^2 \leq (1+\kappa)\|x^+\|_2^2$.

On the event $\mathsf{B}_{\kappa,\alpha,\beta}$ we clearly have $A\mathcal{X}^-, A\mathcal{X}^+ \subset 5\mathbb{B}_2^k$ and $\|Ax^+ - Ax^-\|_2 \geq \frac{\gamma}{2}$ for all $x^+ \in \mathcal{X}^+, x^- \in \mathcal{X}^-$. Further, by applying Lemma 3.6 for $\mathcal{X} = \mathcal{X}^+ - \mathcal{X}^- \subset 2\mathbb{B}_2^d$ and $t = 1 - \varepsilon$, we obtain that on the event $\mathsf{B}_{\kappa,\alpha,\beta}$,

$$\langle \tfrac{Au}{\|Au\|_2}, A(x^+ - x^-) \rangle \geq \sqrt{\tfrac{2}{3}}(1-\varepsilon)\|A(x^+ - x^-)\|_2 - \tfrac{\kappa}{\sqrt{2}} - \sqrt{2}(3 + (1-\varepsilon))\beta - \tfrac{3}{\sqrt{2}}(1+2)\alpha$$

for all $x^+ \in \mathcal{X}^+, x^- \in \mathcal{X}^-$. For $\kappa \lesssim \gamma(1-\varepsilon)$ and $\alpha, \beta \lesssim \gamma(1-\varepsilon)$, we obtain

$$\langle \tfrac{Au}{\|Au\|_2}, A(x^+ - x^-) \rangle \geq \tfrac{1}{2}(1-\varepsilon)\|A(x^+ - x^-)\|_2 \qquad \text{for all } x^+ \in \mathcal{X}^+, x^- \in \mathcal{X}^-.$$

Hence, if $\kappa \lesssim \gamma(1-\varepsilon)$ and $\alpha, \beta \lesssim \gamma(1-\varepsilon)$, then $\mathsf{B}_{\kappa,\alpha,\beta} \subset \mathsf{A}_{\varepsilon',\gamma'}$ for $\varepsilon' = \frac{1+\varepsilon}{2}$ and $\gamma' = \frac{\gamma}{2}$. For this choice of ε' and γ', Corollary 3.4 implies that if $\lambda \geq 5\sqrt{k}$, then

$$\begin{aligned}
&\mathbb{P}_{V',\tau}(H[\sqrt{k}V',\tau] \ c\gamma(1-\varepsilon)\sqrt{k}\text{-separates } A\mathcal{X}^- \text{ from } A\mathcal{X}^+ \mid \mathsf{A}_{\varepsilon',\gamma'})\\
&\geq c\frac{\gamma(1-\varepsilon)}{\lambda}\sqrt{k}\exp(-Ck\log(4(1-\varepsilon)^{-1})).
\end{aligned}$$

Therefore, applying (3.9) with $\rho = c\gamma(1-\varepsilon)\sqrt{k}$ and $\varepsilon' = \frac{1+\varepsilon}{2}, \gamma' = \frac{\gamma}{2}$, we obtain that if $\lambda \geq 5\sqrt{k}$, then

$$\begin{aligned}
&\mathbb{P}(H[g,\tau] \ c\gamma(1-\varepsilon)\sqrt{k}\text{-separates } \mathcal{X}^- \text{ from } \mathcal{X}^+) && (3.10)\\
&\geq c\frac{\gamma(1-\varepsilon)}{\lambda}\sqrt{k}\exp(-Ck\log(4(1-\varepsilon)^{-1})) \cdot \mathbb{P}(\mathsf{A}_{\varepsilon',\gamma'})\\
&\geq c\frac{\gamma(1-\varepsilon)}{\lambda}\sqrt{k}\exp(-Ck\log(4(1-\varepsilon)^{-1})) \cdot \mathbb{P}(\mathsf{B}_{\kappa,\alpha,\beta}),
\end{aligned}$$

where the second inequality holds for $\kappa \lesssim \gamma(1-\varepsilon)$ and $\alpha, \beta \lesssim \gamma(1-\varepsilon)$.

Let $\mathcal{T} \subset \mathbb{R}^d$ be a set. By matrix deviation inequality for Gaussian matrices (see e.g. [Ver18, Section 9.1]), if

$$k \gtrsim \theta^{-2}(w^2(\mathcal{T}) + \log(2/\eta)\,\mathrm{rad}^2(\mathcal{T})),$$

then with probability at least $1 - \eta$,

$$\sup_{x \in \mathcal{T}} \big| \|Ax\|_2 - \|x\|_2 \big| \leq \theta.$$

Hence, a union bound implies that if

$$k \gtrsim \kappa^{-2}\log(2/\eta), \quad k \gtrsim (\alpha^{-2} + \beta^{-2})(w^2(\mathcal{X}^+ - \mathcal{X}^-) + \log(2/\eta)),$$

then $\mathsf{B}_{\kappa,\alpha,\beta}$ occurs with probability at least $1 - \eta$. In particular, if

$$k \gtrsim \gamma^{-2}(1-\varepsilon)^{-2}(w^2(\mathcal{X}^+ - \mathcal{X}^-) + 1),$$

then $\mathsf{B}_{\kappa,\alpha,\beta}$ with $\kappa \sim \gamma(1-\varepsilon)$ and $\alpha, \beta \sim \gamma(1-\varepsilon)$ occurs with probability at least $\frac{1}{2}$. Combining this result with (3.10), we obtain that if

$$k \gtrsim \gamma^{-2}(1-\varepsilon)^{-2}(w^2(\mathcal{X}^+ - \mathcal{X}^-) + 1), \quad \lambda \geq 5\sqrt{k},$$

then

$$\mathbb{P}(H[\boldsymbol{g},\tau]\, c\gamma(1-\varepsilon)\sqrt{k}\text{-separates } \mathcal{X}^- \text{ from } \mathcal{X}^+) \geq c\frac{\gamma(1-\varepsilon)}{2\lambda}\sqrt{k}\exp(-Ck\log(4(1-\varepsilon)^{-1})).$$

Let $\mu := \gamma(1-\varepsilon)$. Setting $k = \mu^{-2}t^2$ completes the proof. $\blacksquare$

3.2 Separation of Two Points

The following result concerns the separation of two arbitrary points by a random hyperplane, and does not follow directly from Theorem 1.6. It can be shown in a more elementary way, leading to a stronger statement.

Theorem 3.7 *There exist absolute constants $c, C > 0$ such that the following holds. Let $x^-, x^+ \in R\mathbb{B}_2^d$. Let $\boldsymbol{g} \in \mathbb{R}^d$ denote a standard Gaussian vector and let $\tau \in [-\lambda, \lambda]$ be uniformly distributed. If $\lambda \geq CR$ then with probability at least $c\|x^+ - x^-\|_2/\lambda$, the hyperplane $H[\boldsymbol{g},\tau]\,\|x^+ - x^-\|_2$-separates x^- from x^+.*

Proof. Since $\frac{x^+ - x^-}{\|x^+ - x^-\|_2} \in \mathbb{S}^{d-1}$, the vector $\langle \boldsymbol{g}, \frac{x^+ - x^-}{\|x^+ - x^-\|_2} \rangle \in \mathbb{S}^{d-1}$ is standard Gaussian. Therefore,

$$\mathbb{P}(\langle \boldsymbol{g}, x^+ - x^- \rangle \geq 4\|x^+ - x^-\|_2) \geq c$$

for an absolute constant $c > 0$. Further, we have the inequalities

$$\mathbb{P}(\langle \boldsymbol{g}, x^+ \rangle \leq \lambda) \geq 1 - \exp(-\lambda^2/2\|x^+\|_2^2) \geq 1 - \exp(-\lambda^2/2R^2)$$

and

$$\mathbb{P}(\langle \boldsymbol{g}, x^- \rangle \geq -\lambda) \geq 1 - \exp(-\lambda^2/2R^2).$$

Define the event

$$\mathsf{A} := \{\langle \boldsymbol{g}, x^+ - x^- \rangle \geq 4\|x^+ - x^-\|_2\} \cap \{\langle \boldsymbol{g}, x^+ \rangle \leq \lambda\} \cap \{\langle \boldsymbol{g}, x^- \rangle \geq -\lambda\}.$$

By the above, $\mathbb{P}(\mathsf{A}) \geq c - 2\exp(-\lambda^2/2R^2)$. Therefore, if $\lambda \geq CR$ for $C > 0$ an absolute constant that is chosen large enough, then $\mathbb{P}(\mathsf{A}) \geq \frac{c}{2}$. Let us show that

$$\mathbb{P}_\tau(H[\boldsymbol{g},\tau]\,\|x^+ - x^-\|_2\text{-separates } x^- \text{ from } x^+ \mid \mathsf{A}) \geq \frac{\|x^+ - x^-\|_2}{\lambda}.$$

Indeed, on the event A it holds $\langle g, x^+ \rangle, \langle g, x^- \rangle \in [-\lambda, \lambda]$ and $\langle g, x^+ - x^- \rangle \geq 4\|x^+ - x^-\|_2$. In particular, the interval $\mathcal{I} := [\langle g, x^- \rangle + \|x^+ - x^-\|_2, \langle g, x^+ \rangle - \|x^+ - x^-\|_2]$ belongs to $[-\lambda, \lambda]$ with $|\mathcal{I}| \geq 2\|x^+ - x^-\|_2$. If $-\tau \in \mathcal{I}$, then $H[g, \tau] \|x^+ - x^-\|_2$-separates x^- from x^+. Therefore,

$$\mathbb{P}_\tau(H[g, \tau] \|x^+ - x^-\|_2\text{-separates } x^- \text{ from } x^+ \mid A) \geq \mathbb{P}_\tau(-\tau \in \mathcal{I} \mid A) = \frac{|\mathcal{I}|}{2\lambda} \geq \frac{\|x^+ - x^-\|_2}{\lambda}.$$

The result now follows from

$$
\begin{aligned}
&\mathbb{P}(H[g, \tau] \|x^+ - x^-\|_2\text{-separates } x^- \text{ from } x^+) \\
&\geq \mathbb{P}(\{H[g, \tau] \|x^+ - x^-\|_2\text{-separates } x^- \text{ from } x^+\} \cap A) \\
&= \mathbb{P}(H[g, \tau] \|x^+ - x^-\|_2\text{-separates } x^- \text{ from } x^+ \mid A) \cdot \mathbb{P}(A) \\
&\geq \frac{\|x^+ - x^-\|_2}{\lambda} \cdot \frac{c}{2}.
\end{aligned}
$$

∎

4 Distance Preservation

The following theorem describes how the Euclidean geometry of two sets $\mathcal{X}^-, \mathcal{X}^+$ is transformed by applying a random ReLU-layer. It shows that, with high probability, Euclidean distances are approximately preserved provided that both the layer is wide and the bias parameter λ is large enough.

Theorem 4.1 *There exist absolute constants $C, C', c > 0$ such that the following holds. Let $\mathcal{X}^-, \mathcal{X}^+ \subset R\mathbb{B}_2^d$ and let $\Phi \colon \mathbb{R}^d \to \mathbb{R}^n$ be a random ReLU-layer with maximal bias $\lambda \geq 0$. If $\lambda \geq CR\sqrt{\log(\lambda^2/\varepsilon)}$, $\lambda^2/\varepsilon \geq e$ and*

$$n \geq C'\varepsilon^{-2}\lambda^2\big(w^2(\mathcal{X}^+) + w^2(\mathcal{X}^-) + u^2\lambda^2\big), \tag{4.1}$$

then with probability at least $1 - 2\exp(-cu^2)$ the following three events occur:

(i) For all $x^+ \in \mathcal{X}^+, x^- \in \mathcal{X}^-$, we have

$$\left| \|\Phi(x^+) - \Phi(x^-)\|_2^2 - \|x^+ - x^-\|_2^2\Big(1 - \sqrt{\tfrac{2}{\pi}} \tfrac{\|x^+ - x^-\|_2}{3\lambda}\Big) \right| \leq \varepsilon. \tag{4.2}$$

(ii) For all $x \in \mathcal{X}^- \cup \mathcal{X}^+$, we have

$$\left| \|\Phi(x)\|_2^2 - \big(\|x\|_2^2 + \tfrac{\lambda^2}{3}\big) \right| \leq \varepsilon. \tag{4.3}$$

(iii) For all $x^+ \in \mathcal{X}^+, x^- \in \mathcal{X}^-$, we have

$$\left| \langle \Phi(x^+), \Phi(x^-) \rangle - \Big(\langle x^+, x^- \rangle + \tfrac{\lambda^2}{3} + \sqrt{\tfrac{2}{\pi}} \tfrac{1}{6} \tfrac{\|x^+ - x^-\|_2^3}{\lambda}\Big) \right| \leq \varepsilon. \tag{4.4}$$

The following two results are straightforward corollaries of (4.3) and (4.2) in Theorem 4.1, respectively.

Corollary 4.2 *There exist absolute constants $C, C' > 0$ such that the following holds. Let $\mathcal{X} \subset R\mathbb{B}_2^d$ and let $\Phi \colon \mathbb{R}^d \to \mathbb{R}^n$ be a random ReLU-layer with maximal bias $\lambda \geq 0$. For any $\eta \in (0, 1)$, if $\lambda \geq CR$ and*

$$n \geq C'\big(\lambda^{-2}w^2(\mathcal{X}) + \log(e/\eta)\big),$$

then $\Phi(\mathcal{X}) \subset \lambda\mathbb{B}_2^n$ with probability at least $1 - \eta$.

Corollary 4.3 *There exist absolute constants $C, C', c > 0$ such that the following holds. Let $\mathcal{X}^-, \mathcal{X}^+ \subset R\mathbb{B}_2^d$ be δ-separated sets and let $\Phi \colon \mathbb{R}^d \to \mathbb{R}^n$ be a random ReLU-layer with maximal bias $\lambda \geq 0$. For any $\eta \in (0, 1)$, if $\lambda \geq CR\sqrt{\log(\lambda/\delta)}$, $\lambda/\delta \geq e$ and*

$$n \geq C'\delta^{-4}\lambda^2\big(w^2(\mathcal{X}^+) + w^2(\mathcal{X}^-) + \log(e/\eta)\lambda^2\big),$$

then $\Phi(\mathcal{X}^-)$ and $\Phi(\mathcal{X}^+)$ are $\frac{\delta}{2}$-separated with probability at least $1 - \eta$.

Let us outline the main steps of the proof of Theorem 4.1. Note that it suffices to show (4.4). Indeed, (4.3) trivially follows from (4.4). Further, (4.2) follows from (4.4) and (4.3) by polarization. To show (4.4), we proceed in two steps:

1. Compute the expected value of $\langle \Phi(x^+), \Phi(x^-) \rangle$ for two arbitrary points $x^+, x^- \in R\mathbb{B}_2^d$ (Proposition 4.4).

2. Show uniform concentration of $\langle \Phi(x^+), \Phi(x^-) \rangle$ around its expected value using a concentration result for empirical product processes due to Mendelson [Men16, Thm. 1.13].

Proposition 4.4 *There exists an absolute constant $C > 0$ such that the following holds. Let $\Phi \colon \mathbb{R}^d \to \mathbb{R}^n$ be a random ReLU-layer with maximal bias $\lambda \geq 0$. If $\lambda \geq CR\sqrt{\log(\lambda^2/\varepsilon)}$ and $\lambda^2/\varepsilon \geq e$, then for any $x^+, x^- \in R\mathbb{B}_2^d$,*

$$\left| \mathbb{E}[\langle \Phi(x^+), \Phi(x^-) \rangle] - \left(\langle x^+, x^- \rangle + \tfrac{\lambda^2}{3} + \sqrt{\tfrac{2}{\pi}} \tfrac{1}{6} \tfrac{\|x^+ - x^-\|_2^3}{\lambda} \right) \right| \leq \varepsilon.$$

To prove Proposition 4.4, we will make use of the following lemma:

Lemma 4.5 *Let $\tau \in [-\lambda, \lambda]$ be uniformly distributed. For $a, b \in [-\lambda, \lambda]$,*

$$\mathbb{E}[\mathrm{ReLU}(a + \tau)\, \mathrm{ReLU}(b + \tau)] = \tfrac{ab}{2} + \tfrac{(\min\{a,b\})^2 \max\{a,b\}}{4\lambda} - \tfrac{(\min\{a,b\})^3}{12\lambda} + (a + b)\tfrac{\lambda}{4} + \tfrac{\lambda^2}{6}.$$

Proof. We may assume that $a \leq b$. Then

$$\begin{aligned}
\mathbb{E}[\mathrm{ReLU}(a + \tau)\, \mathrm{ReLU}(b + \tau)] &= \mathbb{E}[(a + \tau)\mathbb{1}_{\tau \geq -a}(b + \tau)\mathbb{1}_{\tau \geq -b}] \\
&= \mathbb{E}[(a + \tau)(b + \tau)\mathbb{1}_{\tau \geq -a}] \\
&= \tfrac{1}{2\lambda} \int_{-a}^{\lambda} (ab + (a + b)s + s^2)\, ds \\
&= \tfrac{1}{2\lambda} \left(ab(\lambda + a) + (a + b)(\tfrac{\lambda^2}{2} - \tfrac{a^2}{2}) + \tfrac{\lambda^3}{3} + \tfrac{a^3}{3} \right) \\
&= \tfrac{ab}{2} + \tfrac{a^2 b}{4\lambda} - \tfrac{a^3}{12\lambda} + (a + b)\tfrac{\lambda}{4} + \tfrac{\lambda^2}{6} \\
&= \tfrac{ab}{2} + \tfrac{(\min\{a,b\})^2 \max\{a,b\}}{4\lambda} - \tfrac{(\min\{a,b\})^3}{12\lambda} + (a + b)\tfrac{\lambda}{4} + \tfrac{\lambda^2}{6}.
\end{aligned}$$

∎

Proof of Proposition 4.4. Clearly,

$$\mathbb{E}[\langle \Phi(x^+), \Phi(x^-) \rangle] = 2 \cdot \mathbb{E}[\mathrm{ReLU}(\langle g, x^+ \rangle + \tau)\, \mathrm{ReLU}(\langle g, x^- \rangle + \tau)],$$

where g denotes a standard Gaussian vector and $\tau \in [-\lambda, \lambda]$ is an independent and uniformly distributed random variable. Let us begin by showing that

$$\begin{aligned}
\mathbb{E}\bigg[&\tfrac{\langle g, x^+ \rangle \langle g, x^- \rangle}{2} + \tfrac{(\min\{\langle g, x^+ \rangle, \langle g, x^- \rangle\})^2 \max\{\langle g, x^+ \rangle, \langle g, x^- \rangle\}}{4\lambda} - \tfrac{(\min\{\langle g, x^+ \rangle, \langle g, x^- \rangle\})^3}{12\lambda} \\
&+ \tfrac{\lambda(\langle g, x^+ \rangle + \langle g, x^- \rangle)}{4} + \tfrac{\lambda^2}{6} \bigg] \\
&= \tfrac{1}{2} \left(\langle x^+, x^- \rangle + \tfrac{\lambda^2}{3} + \sqrt{\tfrac{2}{\pi}} \tfrac{1}{6} \tfrac{\|x^+ - x^-\|_2^3}{\lambda} \right).
\end{aligned} \tag{4.5}$$

Since $\mathbb{E}[\langle g, x^+ \rangle \langle g, x^- \rangle] = \langle x^+, x^- \rangle$ for any vectors $x^+, x^- \in \mathbb{R}^d$ and $\mathbb{E}[\langle g, x^+ \rangle] = \mathbb{E}[\langle g, x^- \rangle] = 0$, this amounts to showing that

$$\begin{aligned}
\mathbb{E}&\left[3(\min\{\langle g, x^+ \rangle, \langle g, x^- \rangle\})^2 \max\{\langle g, x^+ \rangle, \langle g, x^- \rangle\} - (\min\{\langle g, x^+ \rangle, \langle g, x^- \rangle\})^3 \right] \\
&= \sqrt{\tfrac{2}{\pi}} \|x^+ - x^-\|_2^3.
\end{aligned} \tag{4.6}$$

Since g is symmetric,

$$\mathbb{E}\big[(\min\{\langle g, x^+ \rangle, \langle g, x^- \rangle\})^2 \max\{\langle g, x^+ \rangle, \langle g, x^- \rangle\}\big]$$
$$= -\mathbb{E}\big[(\max\{\langle g, x^+ \rangle, \langle g, x^- \rangle\})^2 \min\{\langle g, x^+ \rangle, \langle g, x^- \rangle\}\big]$$

and

$$\mathbb{E}\big[(\min\{\langle g, x^+ \rangle, \langle g, x^- \rangle\})^3\big] = -\mathbb{E}\big[(\max\{\langle g, x^+ \rangle, \langle g, x^- \rangle\})^3\big].$$

Therefore,

$$\mathbb{E}\Big[3(\min\{\langle g, x^+ \rangle, \langle g, x^- \rangle\})^2 \max\{\langle g, x^+ \rangle, \langle g, x^- \rangle\} - (\min\{\langle g, x^+ \rangle, \langle g, x^- \rangle\})^3\Big]$$
$$= \tfrac{1}{2} \cdot \Big(3\mathbb{E}\big[(\min\{\langle g, x^+ \rangle, \langle g, x^- \rangle\})^2 \max\{\langle g, x^+ \rangle, \langle g, x^- \rangle\}\big]$$
$$\quad - 3\mathbb{E}\big[(\max\{\langle g, x^+ \rangle, \langle g, x^- \rangle\})^2 \min\{\langle g, x^+ \rangle, \langle g, x^- \rangle\}\big]$$
$$\quad - \mathbb{E}\big[(\min\{\langle g, x^+ \rangle, \langle g, x^- \rangle\})^3\big] + \mathbb{E}\big[(\max\{\langle g, x^+ \rangle, \langle g, x^- \rangle\})^3\big] \Big)$$
$$= \tfrac{1}{2} \cdot \mathbb{E}\big[(\max\{\langle g, x^+ \rangle, \langle g, x^- \rangle\} - \min\{\langle g, x^+ \rangle, \langle g, x^- \rangle\})^3\big]$$
$$= \tfrac{1}{2} \cdot \mathbb{E}\big[|\langle g, x^+ - x^- \rangle|^3\big].$$

Using that g is rotation invariant and $\mathbb{E}[|g|^3] = 2\sqrt{\tfrac{2}{\pi}}$ for $g \sim \mathsf{N}(0,1)$, we arrive at (4.6). Define the event

$$\mathsf{A} = \Big\{ \max\{|\langle g, x^+ \rangle|, |\langle g, x^- \rangle|\} \le \lambda \Big\}.$$

Since $\|\langle g, x^+ \rangle\|_{\psi_2}, \|\langle g, x^- \rangle\|_{\psi_2} \lesssim R$, we have

$$\mathbb{P}(\mathsf{A}^C) \le \mathbb{P}(|\langle g, x^+ \rangle| > \lambda) + \mathbb{P}(|\langle g, x^- \rangle| > \lambda) \le 4\exp(-c\lambda^2/R^2)$$

for some absolute constant $c > 0$. By using the Cauchy-Schwarz inequality twice, we obtain

$$\mathbb{E}[|\mathrm{ReLU}(\langle g, x^+ \rangle + \tau)\,\mathrm{ReLU}(\langle g, x^- \rangle + \tau)| \cdot \mathbb{1}_{\mathsf{A}^C}]$$
$$\le \big(\mathbb{E}\big[(\mathrm{ReLU}(\langle g, x^+ \rangle + \tau))^2 (\mathrm{ReLU}(\langle g, x^- \rangle + \tau))^2\big]\big)^{1/2} \cdot \mathbb{P}(\mathsf{A}^C)^{1/2}$$
$$\le \|\mathrm{ReLU}(\langle g, x^+ \rangle + \tau)\|_{L^4} \cdot \|\mathrm{ReLU}(\langle g, x^- \rangle + \tau)\|_{L^4} \cdot 2\exp(-c\lambda^2/2R^2)$$
$$\lesssim \lambda^2 \exp(-c\lambda^2/2R^2)$$
$$\le \varepsilon,$$

where the last inequality follows if $\lambda \ge CR\sqrt{\log(\lambda^2/\varepsilon)}$ for $C > 0$ an absolute constant that is chosen large enough and $\lambda^2/\varepsilon \ge e$. Therefore,

$$\left| \mathbb{E}[\langle \Phi(x^+), \Phi(x^-) \rangle] - \Big(\langle x^+, x^- \rangle + \tfrac{\lambda^2}{3} + \sqrt{\tfrac{2}{\pi}} \tfrac{1}{6} \tfrac{\|x^+ - x^-\|_2^3}{\lambda} \Big) \right|$$
$$= \left| 2\mathbb{E}[\mathrm{ReLU}(\langle g, x^+ \rangle + \tau)\,\mathrm{ReLU}(\langle g, x^- \rangle + \tau)] - \Big(\langle x^+, x^- \rangle + \tfrac{\lambda^2}{3} + \sqrt{\tfrac{2}{\pi}} \tfrac{1}{6} \tfrac{\|x^+ - x^-\|_2^3}{\lambda} \Big) \right|$$
$$\le \left| 2\mathbb{E}[\mathrm{ReLU}(\langle g, x^+ \rangle + \tau)\,\mathrm{ReLU}(\langle g, x^- \rangle + \tau) \cdot \mathbb{1}_{\mathsf{A}}] - \Big(\langle x^+, x^- \rangle + \tfrac{\lambda^2}{3} + \sqrt{\tfrac{2}{\pi}} \tfrac{1}{6} \tfrac{\|x^+ - x^-\|_2^3}{\lambda} \Big) \right|$$
$$+ 2\varepsilon. \tag{4.7}$$

Using Lemma 4.5 and the independence of g and τ we obtain

$$\mathbb{E}[\mathrm{ReLU}(\langle g, x^+ \rangle + \tau)\,\mathrm{ReLU}(\langle g, x^- \rangle + \tau) \cdot \mathbb{1}_{\mathsf{A}}]$$
$$= \mathbb{E}_g[\mathbb{1}_{\mathsf{A}}\mathbb{E}_\tau[\mathrm{ReLU}(\langle g, x^+ \rangle + \tau)\,\mathrm{ReLU}(\langle g, x^- \rangle + \tau)]]$$
$$= \mathbb{E}_g\big[\mathbb{1}_{\mathsf{A}}\big(\tfrac{\langle g, x^+ \rangle \langle g, x^- \rangle}{2} + \tfrac{(\min\{\langle g, x^+ \rangle, \langle g, x^- \rangle\})^2 \max\{\langle g, x^+ \rangle, \langle g, x^- \rangle\}}{4\lambda} - \tfrac{(\min\{\langle g, x^+ \rangle, \langle g, x^- \rangle\})^3}{12\lambda}\big)\big]$$
$$+ \mathbb{E}_g\big[\mathbb{1}_{\mathsf{A}}\big((\langle g, x^+ \rangle + \langle g, x^- \rangle)\tfrac{\lambda}{4} + \tfrac{\lambda^2}{6}\big)\big].$$

In combination with (4.5) and the Cauchy-Schwarz inequality this yields

$$\left| \mathbb{E}[\text{ReLU}(\langle g, x^+\rangle + \tau)\,\text{ReLU}(\langle g, x^-\rangle + \tau) \cdot \mathbb{1}_A] - \tfrac{1}{2}\left(\langle x^+, x^-\rangle + \tfrac{\lambda^2}{3} + \sqrt{\tfrac{2}{\pi}}\tfrac{1}{6}\tfrac{\|x^+ - x^-\|_2^3}{\lambda}\right)\right|$$

$$\leq \mathbb{E}\Big[\Big| \tfrac{\langle g, x^+\rangle\langle g, x^-\rangle}{2} + \tfrac{(\min\{\langle g, x^+\rangle, \langle g, x^-\rangle\})^2 \max\{\langle g, x^+\rangle, \langle g, x^-\rangle\}}{4\lambda} - \tfrac{(\min\{\langle g, x^+\rangle, \langle g, x^-\rangle\})^3}{12\lambda}$$

$$+ (\langle g, x^+\rangle + \langle g, x^-\rangle)\tfrac{\lambda}{4} + \tfrac{\lambda^2}{6}\Big| \cdot \mathbb{1}_{A^C}\Big]$$

$$\lesssim \Big(\|\langle g, x^+\rangle\|_{L^4} \cdot \|\langle g, x^-\rangle\|_{L^4} + \tfrac{1}{\lambda}\big(\|\langle g, x^+\rangle\|_{L^6}^3 + \|\langle g, x^-\rangle\|_{L^6}^3\big)$$

$$+ \lambda\big(\|\langle g, x^+\rangle\|_{L^2} + \|\langle g, x^-\rangle\|_{L^2}\big) + \lambda^2\Big) \cdot \mathbb{P}(A^C)^{1/2}$$

$$\lesssim \lambda^2 \cdot \exp(-c\lambda^2 / 2R^2)$$

$$\leq \varepsilon,$$

where the last two inequalities follow by using $\|\langle g, x^+\rangle\|_{\psi_2}, \|\langle g, x^-\rangle\|_{\psi_2} \lesssim R$, $\lambda \geq CR\sqrt{\log(\lambda^2/\varepsilon)}$ for $C > 0$ an absolute constant that is chosen large enough and $\lambda^2/\varepsilon \geq e$. Together with (4.7) this implies

$$\left| \mathbb{E}[\langle \Phi(x^+), \Phi(x^-)\rangle] - \left(\langle x^+, x^-\rangle + \tfrac{\lambda^2}{3} + \sqrt{\tfrac{2}{\pi}}\tfrac{1}{6}\tfrac{\|x^+ - x^-\|_2^3}{\lambda}\right)\right| \leq 4\varepsilon.$$

By rescaling ε, we obtain the result. ∎

Proof of Theorem 4.1. Let us start by showing (4.4). Let $g \in \mathbb{R}^d$ be a standard Gaussian random vector and $\tau \in [-\lambda, \lambda]$ be uniformly distributed. Since the ReLU is 1-Lipschitz, it follows that

$$\| \text{ReLU}(\langle g, x^+\rangle + \tau) - \text{ReLU}(\langle g, x^-\rangle + \tau)\|_{\psi_2} \lesssim \|\langle g, x^+ - x^-\rangle\|_{\psi_2} \lesssim \|x^+ - x^-\|_2$$

and

$$\| \text{ReLU}(\langle g, x\rangle + \tau)\|_{\psi_2} \lesssim \|x\|_2 + \lambda \leq 2\lambda$$

for all $x^+, x^- \in \mathbb{R}^d$ and $x \in R\mathbb{B}_2^d$. Hence, the stochastic processes

$$\left\{ \text{ReLU}(\langle g, x^+\rangle + \tau)\right\}_{x^+ \in \mathcal{X}^+}, \quad \left\{ \text{ReLU}(\langle g, x^-\rangle + \tau)\right\}_{x^- \in \mathcal{X}^-}$$

are sub-Gaussian with respect to the Euclidean metric and their radii in sub-Gaussian norm are bounded by λ. By a concentration result for empirical product processes where each process is sub-Gaussian [Men16, Thm. 1.13], we have

$$\tfrac{1}{n}\left| \sum_{i=1}^n \text{ReLU}(\langle w_i, x^+\rangle + b_i)\,\text{ReLU}(\langle w_i, x^-\rangle + b_i) - \mathbb{E}[\text{ReLU}(\langle w_i, x^+\rangle + b_i)\,\text{ReLU}(\langle w_i, x^-\rangle + b_i)]\right|$$

$$\leq C \cdot \left(\frac{(w(\mathcal{X}^+) + u\lambda)(w(\mathcal{X}^-) + u\lambda)}{n} + \frac{\lambda(w(\mathcal{X}^+) + w(\mathcal{X}^-)) + u\lambda^2}{\sqrt{n}}\right) \tag{4.8}$$

uniformly for all $x^+ \in \mathcal{X}^+, x^- \in \mathcal{X}^-$ with probability at least $1 - 2\exp(-cu^2)$. The condition on n given by (4.1) now implies that the right hand side of (4.8) is bounded by ε. In combination with Proposition 4.4 this shows (4.4) by using the triangle inequality. Arguing analogously we can show that (4.3) holds. Finally, by polarization, (4.4) and (4.3) imply (4.2). ∎

5 Proof of the Main Result (Theorem 2.3)

The following lemma and especially its Corollary 5.2 are crucial ingredients for the proof of Theorem 2.3. In short, they make the following geometric statement precise: Let $\mathcal{X}^+, \mathcal{X}^- \in \mathbb{R}^d$ be two sets and $\Phi\colon \mathbb{R}^d \to \mathbb{R}^n$ a deterministic ReLU-layer. If for every $x^+ \in \mathcal{X}^+$ there exists at least one "neuron" that separates $\mathcal{X}^-$ from x^+, then the transformations $\Phi(\mathcal{X}^-)$ and $\Phi(\mathcal{X}^+)$ are linearly separable.

Lemma 5.1 *Let* $\mathcal{X}^+, \mathcal{X}^- \subset \mathbb{R}^d$. *For* $W = [w_1, \ldots, w_n]^\mathsf{T} \in \mathbb{R}^{n \times d}$ *and* $b = (b_1, \ldots, b_n) \in \mathbb{R}^n$ *define the associated (deterministic) ReLU-layer* $\Phi \colon \mathbb{R}^d \to \mathbb{R}^n$ *by*

$$\Phi(x) := \sqrt{\tfrac{2}{n}} \cdot \mathrm{ReLU}(Wx + b), \qquad x \in \mathbb{R}^d.$$

Set

$$I := \big\{ i \in [n] \mid \langle w_i, x^- \rangle + b_i \leq 0 \text{ for all } x^- \in \mathcal{X}^- \big\}$$

and for $x^+ \in \mathcal{X}^+$, *define* $I_{x^+}(0) \subset [n]$ *to be the set of all indices* $i \in [n]$ *such that* $H[w_i, b_i]$ *separates* $\mathcal{X}^-$ *from* x^+. *Assume that* $\min_{x^+ \in \mathcal{X}^+} |I_{x^+}(0)| \geq 1$. *Then* $|I| \geq 1$ *and the hyperplane* $H[u, 0]$ *given by the vector* $u \in \mathbb{S}_+^{n-1}$ *with*

$$u_i = \begin{cases} \frac{1}{\sqrt{|I|}}, & i \in I, \\ 0, & \text{otherwise,} \end{cases}$$

separates $\Phi(\mathcal{X}^-)$ *from* $\Phi(\mathcal{X}^+)$. *More precisely,*

$$\langle u, \Phi(x^-) \rangle \leq 0 \qquad \text{for all } x^- \in \mathcal{X}^-,$$
$$\langle u, \Phi(x^+) \rangle \geq \tfrac{1}{n} \sum_{i \in I_{x^+}(0)} |\langle w_i, x^+ \rangle + b_i| \qquad \text{for all } x^+ \in \mathcal{X}^+.$$

Proof. Since $I_{x^+}(0) \subset I$ for every $x^+ \in \mathcal{X}^+$ and $\min_{x^+ \in \mathcal{X}^+} |I_{x^+}(0)| \geq 1$, it follows $|I| \geq 1$. Further, for any $x^- \in \mathcal{X}^-$,

$$\langle u, \Phi(x^-) \rangle = \sum_{i \in I} \frac{1}{\sqrt{|I|}} \cdot \sqrt{\tfrac{2}{n}} \, \mathrm{ReLU}(\underbrace{\langle w_i, x^- \rangle + b_i}_{\leq 0}) = 0.$$

On the other hand, for any $x^+ \in \mathcal{X}^+$,

$$\langle u, \Phi(x^+) \rangle \geq \sum_{i \in I_{x^+}(0)} \frac{1}{\sqrt{|I|}} \cdot \sqrt{\tfrac{2}{n}} \, \mathrm{ReLU}(\langle w_i, x^+ \rangle + b_i) \geq \tfrac{1}{n} \sum_{i \in I_{x^+}(0)} |\langle w_i, x^+ \rangle + b_i|.$$

$$\blacksquare$$

Corollary 5.2 *Let* $\mathcal{X}^+, \mathcal{X}^- \subset \mathbb{R}^d$ *and* $t \geq 0$. *For* $W = [w_1, \ldots, w_n]^\mathsf{T} \in \mathbb{R}^{n \times d}$ *and* $b = (b_1, \ldots, b_n) \in \mathbb{R}^n$ *define the associated (deterministic) ReLU-layer* $\Phi \colon \mathbb{R}^d \to \mathbb{R}^n$ *by*

$$\Phi(x) := \sqrt{\tfrac{2}{n}} \cdot \mathrm{ReLU}(Wx + b), \qquad x \in \mathbb{R}^d.$$

For $x^+ \in \mathcal{X}^+$, *define* $I_{x^+}(t) \subset [n]$ *to be the set of all indices* $i \in [n]$ *such that* $H[w_i, b_i]$ *t-separates* $\mathcal{X}^-$ *from* x^+. *Assume that* $\min_{x^+ \in \mathcal{X}^+} |I_{x^+}(t)| \geq n'$ *for some* $n' \geq 1$. *Then* $\Phi(\mathcal{X}^-)$ *and* $\Phi(\mathcal{X}^+)$ *are linearly separable with margin* $\frac{tn'}{2n}$.

Proof. Set

$$I := \big\{ i \in [n] \mid \langle w_i, x^- \rangle + b_i \leq 0 \text{ for all } x^- \in \mathcal{X}^- \big\}.$$

By Lemma 5.1 we have $|I| \geq 1$ and the hyperplane $H[u, 0]$ given by the vector $u \in \mathbb{S}_+^{n-1}$ with

$$u_i = \begin{cases} \frac{1}{\sqrt{|I|}}, & i \in I, \\ 0, & \text{otherwise,} \end{cases}$$

satisfies

$$\langle u, \Phi(x^-) \rangle \leq 0 \qquad \text{for all } x^- \in \mathcal{X}^-,$$
$$\langle u, \Phi(x^+) \rangle \geq \tfrac{1}{n} \sum_{i \in I_{x^+}(0)} |\langle w_i, x^+ \rangle + b_i| \qquad \text{for all } x^+ \in \mathcal{X}^+,$$

where $I_{x^+}(0)$ is the set of all indices $i \in [n]$ such that $H[w_i, b_i]$ separates $\mathcal{X}^-$ from x^+. Clearly, $I_{x^+}(t) \subset I_{x^+}(0)$, which implies for any $x^+ \in \mathcal{X}^+$,

$$\langle u, \Phi(x^+)\rangle \geq \frac{1}{n} \sum_{i \in I_{x^+}(t)} |\langle w_i, x^+\rangle + b_i| > \frac{tn'}{n}.$$

It follows that the hyperplane $H[u, -\frac{tn'}{2n}]$ separates $\Phi(\mathcal{X}^-)$ and $\Phi(\mathcal{X}^+)$ with margin $\frac{tn'}{2n}$. ∎

In line with our proof sketch in Section 1.4, the following two results describe the effect of the first random ReLU-layer Φ in the setup of Theorem 2.3.

Theorem 5.3 *There exists an absolute constant $c > 0$ such that the following holds. Let $\mathcal{X}^-, \mathcal{X}^+ \subset R\mathbb{B}_2^d$ be δ-separated sets with $N^- := |\mathcal{X}^-|$, $N^+ := |\mathcal{X}^+|$. Let $\Phi: \mathbb{R}^d \to \mathbb{R}^n$ be a random ReLU-layer with maximal bias $\lambda \geq 0$. Suppose that $\lambda \gtrsim R$ and*

$$n \gtrsim \delta^{-1}\lambda \cdot \log(2N^- N^+/\eta). \tag{5.1}$$

Then with probability at least $1 - \eta$, the following event occurs: for every $x^- \in \mathcal{X}^-$ there exists a vector $u_{x^-} \in \mathbb{S}_+^{n-1}$ such that the hyperplane $H[u_{x^-}, 0]$ linearly separates $\Phi(x^-)$ from $\Phi(\mathcal{X}^+)$. The vector u_{x^-} is given by $u_{x^-} = \|u'_{x^-}\|_2^{-1}u'_{x^-}$ for

$$(u'_{x^-})_i = \begin{cases} 1, & (\Phi(x^-))_i = 0, \\ 0, & \text{otherwise,} \end{cases} \tag{5.2}$$

and satisfies

$$\langle u_{x^-}, \Phi(x^-)\rangle \leq 0,$$
$$\langle u_{x^-}, \Phi(x^+)\rangle \geq c\|x^+ - x^-\|_2^2\lambda^{-1} \qquad \text{for all } x^+ \in \mathcal{X}^+.$$

Proof. Let $W = [w_1, \ldots, w_n]^\mathsf{T} \in \mathbb{R}^{n \times d}$ and $b = (b_1, \ldots, b_n) \in \mathbb{R}^n$ be the weight matrix and bias vector of Φ, respectively. For $x^- \in \mathcal{X}^-$, $x^+ \in \mathcal{X}^+$ define $I_{x^-, x^+} \subset [n]$ to be the set of all indices $i \in [n]$ where $H[w_i, b_i]$ separates x^- from x^+ and define the events

$$\mathsf{B}_{x^-, x^+}^i := \{H[w_i, b_i] \mid \|x^+ - x^-\|_2\text{-separates } x^- \text{ from } x^+\}.$$

For $n'(x^-, x^+) \in \{1, \ldots, n\}$ a number that is specified later, set

$$\mathsf{B}_{x^-, x^+, n'(x^-, x^+)} := \left\{ \sum_{i=1}^n \mathbb{1}_{\mathsf{B}_{x^-, x^+}^i} \geq n'(x^-, x^+) \right\}, \quad \mathsf{B}_{x^-} := \bigcap_{x^+ \in \mathcal{X}^+} \mathsf{B}_{x^-, x^+, n'(x^-, x^+)}, \quad \mathsf{B} := \bigcap_{x^- \in \mathcal{X}^-} \mathsf{B}_{x^-}.$$

On the event B, the following holds for every $x^- \in \mathcal{X}^-$: For all $x^+ \in \mathcal{X}^+$ there exist at least $n'(x^-, x^+) \geq 1$ hyperplanes $H[w_i, b_i]$ which $\|x^+ - x^-\|_2$-separate x^- from x^+. By Lemma 5.1, this implies that the following holds on the event B: For every $x^- \in \mathcal{X}^-$ there exists $u_{x^-} \in \mathbb{S}_+^{n-1}$ such that the hyperplane $H[u_{x^-}, 0]$ linearly separates $\Phi(x^-)$ from $\Phi(\mathcal{X}^+)$. More precisely,

$$\langle u_{x^-}, \Phi(x^-)\rangle \leq 0,$$
$$\langle u_{x^-}, \Phi(x^+)\rangle \geq \frac{1}{n} \sum_{i \in I_{x^-, x^+}} |\langle w_i, x^+\rangle + b_i| \geq \frac{n'(x^-, x^+)}{n}\|x^+ - x^-\|_2 \qquad \text{for all } x^+ \in \mathcal{X}^+.$$

Further, Lemma 5.1 shows that $u_{x^-} = \|u'_{x^-}\|_2^{-1}u'_{x^-}$ for u'_{x^-} given by

$$(u'_{x^-})_i = \begin{cases} 1, & \langle w_i, x^-\rangle + b_i \leq 0, \\ 0, & \text{otherwise.} \end{cases}$$

Since $\langle w_i, x^-\rangle + b_i \leq 0$ is equivalent to $(\Phi(x^-))_i = 0$, this shows that u_{x^-} is given as described in (5.2). By the union bound,

$$\mathbb{P}(\mathsf{B}^C) \leq \sum_{x^- \in \mathcal{X}^-, x^+ \in \mathcal{X}^+} \mathbb{P}(\mathsf{B}_{x^-, x^+, n'(x^-, x^+)}^C).$$

Let $i \in [n]$. Theorem 3.7 implies that if $\lambda \gtrsim R$, then $\mathbb{P}(B^i_{x^-,x^+}) \geq c_1 \lambda^{-1} \|x^+ - x^-\|_2$ for some absolute constant $c_1 > 0$. Therefore, the Chernoff bound implies that

$$\mathbb{P}\Big(\sum_{i=1}^n \mathbb{1}_{B^i_{x^-,x^+}} \geq \tfrac{c_1}{2} \lambda^{-1} \|x^+ - x^-\|_2 n \Big) \geq 1 - \exp(-c'\lambda^{-1}\|x^+ - x^-\|_2 n).$$

Setting $n'(x^-, x^+) = \lfloor \tfrac{c_1}{2}\lambda^{-1}\|x^+ - x^-\|_2 n \rfloor$, we obtain

$$\mathbb{P}(B^C_{x^-,x^+,n'(x^-,x^+)}) \leq \exp(-c'\lambda^{-1}\|x^+ - x^-\|_2 n) \leq \exp(-c'\lambda^{-1}\delta n).$$

Hence,

$$\mathbb{P}(B^C) \leq \sum_{x^- \in \mathcal{X}^-, x^+ \in \mathcal{X}^+} \exp(-c'\lambda^{-1}\delta n) \leq \eta,$$

where the last inequality follows from

$$n \gtrsim \delta^{-1}\lambda \cdot \big(\log N^- + \log N^+ + \log(\eta^{-1}) \big).$$

Finally, observe that (5.1) implies that $n'(x^-, x^+) = \lfloor \tfrac{c_1}{2}\lambda^{-1}\|x^+ - x^-\|_2 n \rfloor \geq 1$ for all $x^- \in \mathcal{X}^-, x^+ \in \mathcal{X}^+$. ∎

Theorem 5.4 *There exist absolute constants $c, c' > 0$ such that the following holds. Let $\mathcal{X}^-, \mathcal{X}^+ \subset R\mathbb{B}_2^d$ be δ-separated sets. Let $\lambda > 0$ satisfy $\lambda \gtrsim R\sqrt{\log(\lambda/\delta)}$ and $\lambda/\delta \geq e$. Let $\mathcal{C}^+ = \{c_1^+, \ldots, c_{N^+}^+\} \subset R\mathbb{B}_2^d$ and $\mathcal{C}^- = \{c_1^-, \ldots, c_{N^-}^-\} \subset R\mathbb{B}_2^d$ be δ-separated and form a λ/c'-mutual covering for $\mathcal{X}^+$ and $\mathcal{X}^-$ with components $\mathcal{X}_1^+, \ldots, \mathcal{X}_{N^+}^+ \subset \mathcal{X}^+$ and $\mathcal{X}_1^-, \ldots, \mathcal{X}_{N^-}^- \subset \mathcal{X}^-$.*

Let $\Phi \colon \mathbb{R}^d \to \mathbb{R}^n$ be a random ReLU-layer with maximal bias $\lambda \geq 0$. Suppose that

$$n \gtrsim \lambda^{-2} \cdot \big(w^2(\mathcal{X}^-) + w^2(\mathcal{X}^+)\big) + (\tfrac{\lambda}{\delta})^8 \cdot \log(2N^-N^+/\eta)$$
$$n \gtrsim \lambda^6 \cdot \Big(\max_{l \in [N^-]} \big\{ \mathrm{dist}^{-8}(c_l^-, \mathcal{C}^+) \cdot w^2(\mathcal{X}_l^-) \big\} + \max_{j \in [N^+]} \big\{ \mathrm{dist}^{-8}(c_j^+, \mathcal{C}^-) \cdot w^2(\mathcal{X}_j^+) \big\} \Big). \quad (5.3)$$

Then with probability at least $1 - \eta$, the following two events occur:

1) $\Phi(\mathcal{X}^-), \Phi(\mathcal{X}^+) \subset \lambda\mathbb{B}_2^n$;

2) *For every $l \in [N^-]$ there exists a vector $u_{c_l^-} \in \mathbb{S}_+^{n-1}$ such that*

$$\langle u_{c_l^-}, \Phi(x^-) \rangle - 4c'\lambda^{-1}\mathrm{dist}^2(c_l^-, \mathcal{C}^+) \leq -2c'\lambda^{-1}\mathrm{dist}^2(c_l^-, \mathcal{C}^+) \quad \text{for all } x^- \in \mathcal{X}_l^-,$$
$$\langle u_{c_l^-}, \Phi(x^+) \rangle - 4c'\lambda^{-1}\mathrm{dist}^2(c_l^-, \mathcal{C}^+) \geq 2c\lambda^{-1}\|c_l^- - x^+\|_2^2 \quad \text{for all } x^+ \in \mathcal{X}^+.$$

Further, $\|c_l^- - x^+\|_2 \geq \tfrac{1}{2}\mathrm{dist}(c_l^-, \mathcal{C}^+)$ for every $x^+ \in \mathcal{X}^+$, which implies that the hyperplane

$$H[u_{c_l^-}, -4c'\lambda^{-1}\mathrm{dist}^2(c_l^-, \mathcal{C}^+)]$$

linearly separates $\Phi(\mathcal{X}_l^-)$ from $\Phi(\mathcal{X}^+)$ with margin $\min\{2c', \tfrac{c}{2}\}\lambda^{-1}\mathrm{dist}^2(c_l^-, \mathcal{C}^+)$.

Proof. Since $\mathcal{C}^+ = \{c_1^+, \ldots, c_{N^+}^+\}$ and $\mathcal{C}^- = \{c_1^-, \ldots, c_{N^-}^-\}$ form a λ/c'-mutual covering for $\mathcal{X}^+$ and $\mathcal{X}^-$, there exist $r_1^+, \ldots, r_{N^+}^+ \geq 0$ and $r_1^-, \ldots, r_{N^-}^- \geq 0$ such that

1. the sets $\mathcal{X}_j^+ := \mathcal{X}^+ \cap \mathbb{B}_2^d(c_j^+, r_j^+)$ for $j \in [N^+]$, and $\mathcal{X}_l^- := \mathcal{X}^- \cap \mathbb{B}_2^d(c_l^-, r_l^-)$ for $l \in [N^-]$, cover $\mathcal{X}^+$ and $\mathcal{X}^-$, respectively;

2. $r_j^+ \leq c'\lambda^{-1}\mathrm{dist}^2(c_j^+, \mathcal{C}^-)$ for $j \in [N^+]$, and $r_l^- \leq c'\lambda^{-1}\mathrm{dist}^2(c_l^-, \mathcal{C}^+)$ for $l \in [N^-]$.

By Corollary 4.2, if

$$\lambda \gtrsim R, \quad n \gtrsim \lambda^{-2}\big(w^2(\mathcal{X}^-) + w^2(\mathcal{X}^+)\big) + \log(e/\eta),$$

then $\Phi(\mathcal{X}^-), \Phi(\mathcal{X}^+) \subset \lambda\mathbb{B}_2^n$ with probability at least $1 - \eta$. Define A to be the event where for every $l \in [N^-]$ there exists a vector $\boldsymbol{u}_{\boldsymbol{c}_l^-} \in \mathbb{S}_+^{n-1}$ such that

$$\langle \boldsymbol{u}_{\boldsymbol{c}_l^-}, \Phi(\boldsymbol{c}_l^-) \rangle \le 0,$$
$$\langle \boldsymbol{u}_{\boldsymbol{c}_l^-}, \Phi(\boldsymbol{c}_j^+) \rangle \ge c\lambda^{-1}\|\boldsymbol{c}_j^+ - \boldsymbol{c}_l^-\|_2^2 \quad \text{for all } j \in [N^+].$$

Applying Theorem 5.3 to $\mathcal{C}^+$ and $\mathcal{C}^-$, the condition (5.3) implies $\mathbb{P}(A) \ge 1 - \eta$. Define B to be the event where:

(i) For all $l \in [N^-]$:

$$\sup_{\boldsymbol{x}^- \in \mathcal{X}_l^-} \left| \|\Phi(\boldsymbol{x}^-) - \Phi(\boldsymbol{c}_l^-)\|_2^2 - \|\boldsymbol{x}^- - \boldsymbol{c}_l^-\|_2^2 \Big(1 - \sqrt{\tfrac{2}{\pi}} \tfrac{\|\boldsymbol{x}^- - \boldsymbol{c}_l^-\|_2}{3\lambda}\Big) \right| \le \big(c'\lambda^{-1}\operatorname{dist}^2(\boldsymbol{c}_l^-, \mathcal{C}^+)\big)^2,$$

(ii) For all $j \in [N^+]$:

$$\sup_{\boldsymbol{x}^+ \in \mathcal{X}_j^+} \left| \|\Phi(\boldsymbol{x}^+) - \Phi(\boldsymbol{c}_j^+)\|_2^2 - \|\boldsymbol{x}^+ - \boldsymbol{c}_j^+\|_2^2 \Big(1 - \sqrt{\tfrac{2}{\pi}} \tfrac{\|\boldsymbol{x}^+ - \boldsymbol{c}_j^+\|_2}{3\lambda}\Big) \right| \le \big(c'\lambda^{-1}\operatorname{dist}^2(\boldsymbol{c}_j^+, \mathcal{C}^-)\big)^2.$$

By Theorem 4.1 and a union bound, the condition (5.3) implies $\mathbb{P}(B) \ge 1 - \eta$. Let us show that on the event $A \cap B$ the second event from Theorem 5.4 holds. Let $l \in [N^-]$. For any $\boldsymbol{x}^- \in \mathcal{X}_l^-$,

$$\begin{aligned}
\langle \boldsymbol{u}_{\boldsymbol{c}_l^-}, \Phi(\boldsymbol{x}^-) \rangle &= \langle \boldsymbol{u}_{\boldsymbol{c}_l^-}, \Phi(\boldsymbol{c}_l^-) \rangle + \langle \boldsymbol{u}_{\boldsymbol{c}_l^-}, \Phi(\boldsymbol{x}^-) - \Phi(\boldsymbol{c}_l^-) \rangle \\
&\le \|\Phi(\boldsymbol{x}^-) - \Phi(\boldsymbol{c}_l^-)\|_2 \\
&\le \|\boldsymbol{x}^- - \boldsymbol{c}_l^-\|_2 + c'\lambda^{-1}\operatorname{dist}^2(\boldsymbol{c}_l^-, \mathcal{C}^+) \\
&\le r_l^- + c'\lambda^{-1}\operatorname{dist}^2(\boldsymbol{c}_l^-, \mathcal{C}^+) \\
&\le 2c'\lambda^{-1}\operatorname{dist}^2(\boldsymbol{c}_l^-, \mathcal{C}^+).
\end{aligned} \tag{5.4}$$

Let $\boldsymbol{x}^+ \in \mathcal{X}^+$. Then there exists $j \in [N^+]$ such that $\boldsymbol{x}^+ \in \mathcal{X}_j^+$. It holds

$$\begin{aligned}
\langle \boldsymbol{u}_{\boldsymbol{c}_l^-}, \Phi(\boldsymbol{x}^+) \rangle &= \langle \boldsymbol{u}_{\boldsymbol{c}_l^-}, \Phi(\boldsymbol{c}_j^+) \rangle + \langle \boldsymbol{u}_{\boldsymbol{c}_l^-}, \Phi(\boldsymbol{x}^+) - \Phi(\boldsymbol{c}_j^+) \rangle \\
&\ge c\|\boldsymbol{c}_j^+ - \boldsymbol{c}_l^-\|_2^2 \lambda^{-1} - \|\Phi(\boldsymbol{x}^+) - \Phi(\boldsymbol{c}_j^+)\|_2 \\
&\ge c\|\boldsymbol{c}_j^+ - \boldsymbol{c}_l^-\|_2^2 \lambda^{-1} - \|\boldsymbol{x}^+ - \boldsymbol{c}_j^+\|_2 - c'\lambda^{-1}\operatorname{dist}^2(\boldsymbol{c}_j^+, \mathcal{C}^-) \\
&\ge c\|\boldsymbol{c}_j^+ - \boldsymbol{c}_l^-\|_2^2 \lambda^{-1} - 2c'\lambda^{-1}\operatorname{dist}^2(\boldsymbol{c}_j^+, \mathcal{C}^-) \\
&\ge c\|\boldsymbol{c}_j^+ - \boldsymbol{c}_l^-\|_2^2 \lambda^{-1} - 2c'\lambda^{-1}\|\boldsymbol{c}_j^+ - \boldsymbol{c}_l^-\|_2^2 \\
&\ge \tfrac{c}{2}\|\boldsymbol{c}_j^+ - \boldsymbol{c}_l^-\|_2^2 \lambda^{-1},
\end{aligned}$$

where the last inequality follows if $c' \le \tfrac{c}{4}$. If $\lambda \gtrsim R$ and $c' \le 1$, then

$$\begin{aligned}
\|\boldsymbol{x}^+ - \boldsymbol{c}_l^-\|_2 &\le \|\boldsymbol{x}^+ - \boldsymbol{c}_j^+\|_2 + \|\boldsymbol{c}_j^+ - \boldsymbol{c}_l^-\|_2 \\
&\le c'\lambda^{-1}\operatorname{dist}^2(\boldsymbol{c}_j^+, \mathcal{C}^-) + \|\boldsymbol{c}_j^+ - \boldsymbol{c}_l^-\|_2 \\
&\le 2\|\boldsymbol{c}_j^+ - \boldsymbol{c}_l^-\|_2.
\end{aligned}$$

Therefore, for any $\boldsymbol{x}^+ \in \mathcal{X}^+$,

$$\langle \boldsymbol{u}_{\boldsymbol{c}_l^-}, \Phi(\boldsymbol{x}^+) \rangle \ge \tfrac{c}{8}\|\boldsymbol{x}^+ - \boldsymbol{c}_l^-\|_2^2 \lambda^{-1}. \tag{5.5}$$

Subtracting $4c'\lambda^{-1}\operatorname{dist}^2(c_l^-,\mathcal{C}^+)$ in (5.4) and (5.5), we obtain that for all $x^- \in \mathcal{X}_l^-$,

$$\langle u_{c_l^-}, \Phi(x^-)\rangle - 4c'\lambda^{-1}\operatorname{dist}^2(c_l^-,\mathcal{C}^+) \leq -2c'\lambda^{-1}\operatorname{dist}^2(c_l^-,\mathcal{C}^+)$$

and for all $x^+ \in \mathcal{X}^+$,

$$\langle u_{c_l^-}, \Phi(x^+)\rangle - 4c'\lambda^{-1}\operatorname{dist}^2(c_l^-,\mathcal{C}^+) \geq \tfrac{c}{8}\|x^+ - c_l^-\|_2^2\lambda^{-1} - 4c'\lambda^{-1}\operatorname{dist}^2(c_l^-,\mathcal{C}^+).$$

If $\lambda \gtrsim R$ and $c' \leq 1$, then for any $x^+ \in \mathcal{X}_j^+$,

$$
\begin{aligned}
\|c_j^+ - c_l^-\|_2 &\leq \|c_j^+ - x^+\|_2 + \|x^+ - c_l^-\|_2 \\
&\leq c'\lambda^{-1}\operatorname{dist}^2(c_j^+,\mathcal{C}^-) + \|x^+ - c_l^-\|_2 \\
&\leq \tfrac{1}{2}\|c_j^+ - c_l^-\|_2 + \|x^+ - c_l^-\|_2,
\end{aligned}
$$

which implies $\|c_j^+ - c_l^-\|_2 \leq 2\|x^+ - c_l^-\|_2$. In particular, $\|x^+ - c_l^-\|_2 \geq \tfrac{1}{2}\operatorname{dist}(c_l^-,\mathcal{C}^+)$ for all $x^+ \in \mathcal{X}^+$. Furthermore,

$$4c'\lambda^{-1}\operatorname{dist}^2(c_l^-,\mathcal{C}^+) \leq 4c'\lambda^{-1}\|c_l^- - c_j^+\|_2^2 \leq 16c'\lambda^{-1}\|x^+ - c_l^-\|_2^2$$

for any $x^+ \in \mathcal{X}_j^+$. Hence, for every $x^+ \in \mathcal{X}^+$,

$$
\begin{aligned}
\langle u_{c_l^-}, \Phi(x^+)\rangle - 4c'\lambda^{-1}\operatorname{dist}^2(c_l^-,\mathcal{C}^+) &\geq \tfrac{c}{8}\|x^+ - c_l^-\|_2^2\lambda^{-1} - 16c'\lambda^{-1}\|x^+ - c_l^-\|_2^2 \\
&\geq \tfrac{c}{16}\lambda^{-1}\|x^+ - c_l^-\|_2^2,
\end{aligned}
$$

where the last inequality follows if $c' \leq \tfrac{c}{256}$. ∎

The final ingredient for the proof of Theorem 2.3 is the following lemma. It provides a sufficient condition under which we have that $w(\Phi(\mathcal{X})) \lesssim w(\mathcal{X})$ with high probability for $\mathcal{X} \subset \mathbb{R}^d$ and $\Phi\colon \mathbb{R}^d \to \mathbb{R}^n$ a random ReLU-layer.

Lemma 5.5 *Let $\mathcal{X} \subset \mathbb{R}^d$ and $\Phi\colon \mathbb{R}^d \to \mathbb{R}^n$ be a random ReLU-layer with standard Gaussian weight matrix $W \in \mathbb{R}^{n\times d}$ and maximal bias $\lambda \geq 0$. Then $w(\Phi(\mathcal{X})) \leq w(\sqrt{\tfrac{2}{n}}W\mathcal{X})$. Further,*

  i) *If $n \gtrsim \log(2/\eta)$, then $w(\tfrac{1}{\sqrt{n}}W\mathcal{X}) \lesssim w(\mathcal{X}) + \sqrt{n}\operatorname{diam}(\mathcal{X})$ with probability at least $1 - \eta$.*

  ii) *If $n \gtrsim w^2(\operatorname{cone}(\mathcal{X} - \mathcal{X}) \cap \mathbb{S}^{d-1}) + \log(2/\eta)$, then with probability at least $1 - \eta$,*

  $$\sup_{x\neq x'\in\mathcal{X}} \|\tfrac{1}{\sqrt{n}}W(\tfrac{x-x'}{\|x-x'\|_2})\|_2 \leq 2.$$

  On this event $w(\tfrac{1}{\sqrt{n}}W\mathcal{X}') \leq 2w(\mathcal{X}')$ and therefore $w(\Phi(\mathcal{X}')) \leq 2^{3/2}w(\mathcal{X}')$ for every $\mathcal{X}' \subset \mathcal{X}$.

Proof. Let us write $\Phi(x) = \operatorname{ReLU}(T(x))$ for

$$T(x) = \sqrt{\tfrac{2}{n}}Wx + \sqrt{\tfrac{2}{n}}b,$$

where $W \in \mathbb{R}^{n\times d}$ is a standard Gaussian random matrix and b is uniformly distributed on $[-\lambda,\lambda]^n$. Since ReLU is 1-Lipschitz, $w(\Phi(\mathcal{X})) \leq w(T(\mathcal{X}))$ by the Sudakov-Fernique inequality (see e.g. [Ver18, Theorem 7.2.11]). Let $g \in \mathbb{R}^n$ denote a standard Gaussian vector. Since $\mathbb{E}_g\langle g, x\rangle = 0$ for every vector x, it follows

$$
\begin{aligned}
w(T(\mathcal{X})) &= \mathbb{E}_g\Big[\sup_{x\in\mathcal{X}}\langle g, \sqrt{\tfrac{2}{n}}Wx + \sqrt{\tfrac{2}{n}}b\rangle\Big] \\
&= \mathbb{E}_g\Big[\sup_{x\in\mathcal{X}}\langle g, \sqrt{\tfrac{2}{n}}Wx\rangle\Big].
\end{aligned}
$$

Therefore, $w(\Phi(\mathcal{X})) \leq w(\sqrt{\frac{2}{n}} W \mathcal{X})$. Since $w(\mathcal{S}) \leq \frac{\sqrt{n}}{2} \operatorname{diam}(\mathcal{S})$ for any $\mathcal{S} \subset \mathbb{R}^n$, it follows $w(\frac{1}{\sqrt{n}} W \mathcal{X}) \leq \frac{1}{2} \operatorname{diam}(W \mathcal{X})$. By Gaussian projection (see e.g. [Ver18, Section 7.7]), there exists an absolute constant $C > 0$ such that

$$\operatorname{diam}(W\mathcal{X}) \leq C(w(\mathcal{X}) + \sqrt{n}\operatorname{diam}(\mathcal{X}))$$

with probability at least $1 - 2\exp(-n)$. Define

$$\|\tfrac{1}{\sqrt{n}}W\|_{\mathcal{X}} := \sup_{x \neq x' \in \mathcal{X}} \|\tfrac{1}{\sqrt{n}}W(\tfrac{x-x'}{\|x-x'\|_2})\|_2.$$

Let $\mathcal{X}' \subset \mathcal{X}$. Then $\|\frac{1}{\sqrt{n}}Wx - \frac{1}{\sqrt{n}}Wx'\|_2 \leq \|\frac{1}{\sqrt{n}}W\|_{\mathcal{X}}\|x - x'\|_2$ for all $x, x' \in \mathcal{X}'$, which implies $w(\frac{1}{\sqrt{n}}W\mathcal{X}') \leq \|\frac{1}{\sqrt{n}}W\|_{\mathcal{X}}w(\mathcal{X}')$ by the Sudakov-Fernique inequality. By a Gaussian deviation inequality (see e.g. [Ver18, Section 9.1]), if

$$n \gtrsim w^2(\operatorname{cone}(\mathcal{X} - \mathcal{X}) \cap \mathbb{S}^{d-1}) + \log(2/\eta),$$

then $\|\frac{1}{\sqrt{n}}W\|_{\mathcal{X}} \leq 2$ with probability at least $1 - \eta$. ∎

We are ready to prove the main result of this work:

Proof of Theorem 2.3. Let $W = [w_1, \ldots, w_n]^\mathsf{T} \in \mathbb{R}^{n \times d}$ and $b = (b_1, \ldots, b_n) \in [-\lambda, \lambda]^n$ be the Gaussian weight matrix and bias vector of the random ReLU-layer Φ, and let $\hat{W} = [\hat{w}_1, \ldots, \hat{w}_{\hat{n}}]^\mathsf{T} \in \mathbb{R}^{\hat{n} \times n}$ and $\hat{b} = (\hat{b}_1, \ldots, \hat{b}_{\hat{n}}) \in [-\hat{\lambda}, \hat{\lambda}]^{\hat{n}}$ be the Gaussian weight matrix and bias vector of the random ReLU-layer $\hat{\Phi}$. Since $\mathcal{X}^+$ and $\mathcal{X}^-$ have $(R, \delta, C'\lambda)$-mutual complexity (N^+, N^-, w^+, w^-), there exists a $C'\lambda$-mutual covering $\mathcal{C}^+ = \{c_1^+, \ldots, c_{N^+}^+\} \subset \mathbb{R}^d$ and $\mathcal{C}^- = \{c_1^-, \ldots, c_{N^-}^-\} \subset \mathbb{R}^d$ for $\mathcal{X}^+$ and $\mathcal{X}^-$ such that

i) $\max_{j \in [N^+]} w(\mathcal{X}_j^+) \leq w^+$ and $\max_{l \in [N^-]} w(\mathcal{X}_l^-) \leq w^-$;

ii) $\mathcal{C}^+, \mathcal{C}^- \subset R\mathbb{B}_2^d$ are δ-separated.

Here, $\mathcal{X}_1^+, \ldots, \mathcal{X}_{N^+}^+ \subset \mathcal{X}^+$ and $\mathcal{X}_1^-, \ldots, \mathcal{X}_{N^-}^- \subset \mathcal{X}^-$ are the components of the covering. Let $C' := \frac{1}{c'}$, where $c' > 0$ is the absolute constant from Theorem 5.4. By Theorem 5.4, condition (2.1) implies that with probability at least $1 - \eta$, the following event A occurs:

1) $\Phi(\mathcal{X}^-), \Phi(\mathcal{X}^+) \subset \lambda\mathbb{B}_2^n$;

2) For every $l \in [N^-]$ there exists a vector $u_{c_l^-} \in \mathbb{S}_+^{n-1}$ such that the hyperplane

$$H[u_{c_l^-}, -4c'\lambda^{-1}\operatorname{dist}^2(c_l^-, \mathcal{C}^+)]$$

linearly separates $\Phi(\mathcal{X}_l^-)$ from $\Phi(\mathcal{X}^+)$ with margin $\min\{2c', \frac{c}{2}\}\lambda^{-1}\operatorname{dist}^2(c_l^-, \mathcal{C}^+)$.

Here, $c > 0$ is the absolute constant from Theorem 5.4. By Proposition 3.2 (iv), on the event A, for every $l \in [N^-]$ the sets $\Phi(\mathcal{X}_l^-)$ and $\Phi(\mathcal{X}^+)$ are contained in $\lambda\mathbb{B}_2^n$ and are $(\varepsilon_l, \gamma_l)$-linearly separable with

$$\varepsilon_l = 1 - \min\{2c', \tfrac{c}{2}\}\lambda^{-2}\operatorname{dist}^2(c_l^-, \mathcal{C}^+), \quad \gamma_l = \min\{4c', c\}\lambda^{-1}\operatorname{dist}^2(c_l^-, \mathcal{C}^+).$$

For $l \in [N^-]$, $t_l \gtrsim w(\Phi(\mathcal{X}_l^-) - \Phi(\mathcal{X}^+)) + \lambda$ and $i \in [\hat{n}]$, define the event

$$\mathrm{B}_l^i(t_l) := \{H[\hat{w}_i, \hat{b}_i] \; t_l\text{-separates } \Phi(\mathcal{X}^+) \text{ from } \Phi(\mathcal{X}_l^-)\}.$$

Set $\mu_l = \gamma_l(1 - \varepsilon_l)$. By Theorem 1.6, if $\hat{\lambda} \gtrsim \lambda t_l \mu_l^{-1}$, then $\mathbb{P}(\mathrm{B}_l^i(t_l) \mid \mathrm{A}) \geq p_l$ for

$$p_l = \tfrac{t_l}{\hat{\lambda}}\exp(-Ct_l^2\mu_l^{-2}\log(4(1 - \varepsilon_l)^{-1})).$$

Define A′ to be the event where

$$\sup_{x \neq x' \in \mathcal{X}^-} \|\tfrac{1}{\sqrt{n}} W (\tfrac{x-x'}{\|x-x'\|_2})\|_2 \leq 2, \qquad \sup_{x \neq x' \in \mathcal{X}^+} \|\tfrac{1}{\sqrt{n}} W (\tfrac{x-x'}{\|x-x'\|_2})\|_2 \leq 2.$$

By Lemma 5.5 and the union bound, condition (2.1) implies $\mathbb{P}(A') \geq 1 - \eta$. Further, Lemma 5.5 shows that on A′,

$$w(\Phi(\mathcal{X}_l^-) - \Phi(\mathcal{X}^+)) = w(\Phi(\mathcal{X}_l^-)) + w(\Phi(\mathcal{X}^+)) \leq 2^{3/2}(w(\mathcal{X}_l^-) + w(\mathcal{X}^+))$$

for every $l \in [N^-]$. Using that $\mu_l \gtrsim \lambda^{-3}\delta^4$ and $1 - \varepsilon_l \gtrsim \lambda^{-2}\delta^2$ for every $l \in [N^-]$, we obtain that for every $i \in [\hat{n}]$, $l \in [N^-]$ and $t \asymp w^- + w(\mathcal{X}^+) + \lambda$, if $\hat{\lambda} \gtrsim \lambda^4\delta^{-4}t$, then $\mathbb{P}(B_l^i(t) \mid A \cap A') \geq p(t)$ for

$$p(t) = \tfrac{t}{\hat{\lambda}} \exp(-Ct^2\lambda^6\delta^{-8}\log(\lambda/\delta)).$$

Define the events

$$B_l(t) := \Big\{ \sum_{i=1}^{\hat{n}} \mathbb{1}_{B_l^i(t)} \geq \tfrac{p(t)}{2} \hat{n} \Big\}, \quad B_t := \bigcap_{l \in [N^-]} B_l(t).$$

By Chernoff's inequality there exists an absolute constant $c > 0$ such that for all $l \in [N^-]$,

$$\mathbb{P}(B_l(t) \mid A \cap A') \geq 1 - \exp(-cp(t)\hat{n}).$$

On the event B_t, for every $l \in [N^-]$ at least $\tfrac{p(t)}{2}\hat{n}$ out of the $\hat{n}$ hyperplanes $H[\hat{w}_i, \hat{b}_i]$ t-separate $\Phi(\mathcal{X}^+)$ from $\Phi(\mathcal{X}_l^-)$. Since

$$\Phi(\mathcal{X}^-) = \bigcup_{l \in [N^-]} \Phi(\mathcal{X}_l^-),$$

Corollary 5.2 implies that $F(\mathcal{X}^+)$ and $F(\mathcal{X}^-)$ are linearly separable with margin

$$\tfrac{tp(t)}{4} \asymp \tfrac{(w^- + w(\mathcal{X}^+) + \lambda)^2}{\hat{\lambda}} \cdot \exp(-C(w^- + w(\mathcal{X}^+) + \lambda)^2\lambda^6\delta^{-8}\log(\lambda/\delta)).$$

Define B′ to be the event where $F(\mathcal{X}^-), F(\mathcal{X}^+) \subset \hat{\lambda}\mathbb{B}_2^{\hat{n}}$. On the event $B_t \cap B'$ the conclusion of Theorem 2.3 holds. We have

$$\begin{aligned}
\mathbb{P}(B_t \cap B') &\geq \mathbb{P}(B_t \cap B' \cap A \cap A') \\
&= \mathbb{P}(B_t \cap B' \mid A \cap A') \cdot \mathbb{P}(A \cap A') \\
&\geq (1 - \mathbb{P}(B_t^C \mid A \cap A') - \mathbb{P}((B')^C \mid A \cap A')) \cdot (1 - \mathbb{P}(A^C) - \mathbb{P}((A')^C)) \\
&\geq (1 - \mathbb{P}(B_t^C \mid A \cap A') - \mathbb{P}((B')^C \mid A \cap A')) \cdot (1 - 2\eta).
\end{aligned}$$

The union bound implies

$$\mathbb{P}(B_t^C \mid A \cap A') \leq \sum_{l \in [N^-]} \mathbb{P}((B_l(t))^C \mid A \cap A') \leq \exp(\log(N^-) - cp(t)\hat{n}) \leq \eta,$$

where the last inequality follows from

$$\hat{n} \gtrsim (p(t))^{-1}\log(N^-/\eta).$$

On the event $A \cap A'$ it holds $\Phi(\mathcal{X}^-), \Phi(\mathcal{X}^+) \subset \lambda\mathbb{B}_2^n$ and $w(\Phi(\mathcal{X}^-)) \leq 2^{3/2}w(\mathcal{X}^-), w(\Phi(\mathcal{X}^+)) \leq 2^{3/2}w(\mathcal{X}^+)$. Consequently, by Corollary 4.2, if $\hat{\lambda} \gtrsim \lambda$ and

$$\hat{n} \gtrsim (\hat{\lambda})^{-2}(w^2(\mathcal{X}^+) + w^2(\mathcal{X}^-)) + \log(e/\eta), \tag{5.6}$$

then $\mathbb{P}((B')^C \mid A \cap A') \leq \eta$. By Lemma 6.1,

$$w(\mathcal{X}^-) = w(\bigcup_{l \in [N^-]} \mathcal{X}_l^-) \leq w^- + CR\sqrt{\log N^-}$$

for $C > 0$ an absolute constant. Therefore, condition (2.2) implies $\hat{\lambda} \gtrsim \lambda$ and (5.6). In total, we have $\mathbb{P}(B_t \cap B') \geq (1 - 2\eta)^2 \geq 1 - 4\eta$. This completes the proof. ∎

6 Proofs of Special-Case Results

To apply our main result Theorem 2.3 to various special cases, the following lemma will prove to be very useful. Although the inequalities stated therein are well known (for the first inequality see e.g. [JC17, Lemma 10]), we give a proof for the sake of completeness.

Lemma 6.1 *There exists an absolute constant $C > 0$ such that the following holds. Let $\mathcal{X}_j \subset R\mathbb{B}_2^d$ for $j \in [N]$. Then*

$$w\big(\bigcup_{j\in[N]} \mathcal{X}_j\big) \leq \max_{j\in[N]} w(\mathcal{X}_j) + C \cdot R\sqrt{\log N}. \tag{6.1}$$

If all sets $\mathcal{X}_j$ additionally satisfy $\mathrm{diam}(\mathcal{X}_j) \leq r$, then

$$w\big(\bigcup_{j\in[N]} \mathcal{X}_j\big) \leq C \cdot \big(r\sqrt{d} + R\sqrt{\log N}\big). \tag{6.2}$$

Further, if all sets $\mathcal{X}_j$ are additionally finite with $\mathrm{diam}(\mathcal{X}_j) \leq r_j$, then

$$w\big(\bigcup_{j\in[N]} \mathcal{X}_j\big) \leq C \cdot \big(\max_{j\in[N]} (r_j\sqrt{\log|\mathcal{X}_j|}) + R\sqrt{\log N}\big). \tag{6.3}$$

Proof. Let us start by showing (6.1). Let $g \in \mathbb{R}^d$ denote a standard Gaussian random vector and for $j \in [N]$ pick $c_j \in \mathcal{X}_j$. Then $\mathrm{rad}(\mathcal{X}_j - c_j) \leq 2R$. Set

$$X_j := \sup_{x\in\mathcal{X}_j} \langle g, x - c_j\rangle.$$

Then

$$w\big(\bigcup_{j\in[N]} \mathcal{X}_j\big) = \mathbb{E}\max_{j\in[N]}(X_j + \langle g, c_j\rangle)$$

$$\leq \max_{j\in[N]} \mathbb{E}X_j + \mathbb{E}\max_{j\in[N]}(X_j - \mathbb{E}X_j) + \mathbb{E}\max_{j\in[N]}\langle g, c_j\rangle.$$

Clearly, $\mathbb{E}X_j = w(\mathcal{X}_j)$. By Gaussian Lipschitz concentration (see e.g. [FR13, Theorem 8.34]), $X_j - \mathbb{E}X_j$ is a sub-Gaussian random variable with $\|X_j - \mathbb{E}X_j\|_{\psi_2} \lesssim \mathrm{rad}(\mathcal{X}_j - c_j) \leq 2R$. Further, the random variables $\langle g, c_j\rangle$ are sub-Gaussian with $\|\langle g, c_j\rangle\|_{\psi_2} \lesssim \|c_j\|_2 \leq R$. Inequality (6.1) now follows by applying the maximal inequality for sub-Gaussian random variables (see e.g. [BLM13, Theorem 2.5]). Inequalities (6.2) and (6.3) immediately follow from (6.1) by using the standard estimates $w(\mathbb{B}_2^d) \lesssim \sqrt{d}$ and $w(\mathcal{X}) \lesssim \sqrt{\log(|\mathcal{X}|)}$ for any finite $\mathcal{X} \subset \mathbb{B}_2^d$. ∎

6.1 Proof of Theorem 1.3

Define $\mathcal{C}^+ := \mathcal{X}^+$ and $\mathcal{C}^- := \mathcal{X}^-$. Then $\mathcal{C}^+, \mathcal{C}^- \subset \mathbb{B}_2^d$ are δ-separated. We may write $\mathcal{C}^+ = \{c_1^+, \ldots, c_{N^+}^+\}$ and $\mathcal{C}^- = \{c_1^-, \ldots, c_{N^-}^-\}$. Clearly, the sets $\mathcal{X}_j^+ := \mathcal{X}^+ \cap \mathbb{B}_2^d(c_j^+, 0) = \{c_j^+\}$ for $j \in [N^+]$ and $\mathcal{X}_l^- := \mathcal{X}^- \cap \mathbb{B}_2^d(c_l^-, 0) = \{c_l^-\}$ for $l \in [N^-]$ cover $\mathcal{X}^+$ and $\mathcal{X}^-$, respectively. Let $C' > 0$ denote the absolute constant from Theorem 2.3. Then

$$0 \leq \tfrac{1}{C'\lambda}\mathrm{dist}^2(c_j^+, \mathcal{C}^-), \quad 0 \leq \tfrac{1}{C'\lambda}\mathrm{dist}^2(c_l^-, \mathcal{C}^+)$$

for all $j \in [N^+]$ and $l \in [N^-]$. Therefore, $\mathcal{C}^+$ and $\mathcal{C}^-$ form a $C'\lambda$-mutual covering for $\mathcal{X}^+$ and $\mathcal{X}^-$. Moreover, $\mathcal{X}^+$ and $\mathcal{X}^-$ have $(1, \delta, C'\lambda)$-mutual complexity (N^+, N^-, w^+, w^-) with $w^+ = w^- = 0$. Using $w(\mathcal{X}^+) \lesssim \sqrt{\log N^+}$ and

$$w^2(\mathrm{cone}(\mathcal{X}^- - \mathcal{X}^-) \cap \mathbb{S}^{d-1}) + w^2(\mathrm{cone}(\mathcal{X}^+ - \mathcal{X}^+) \cap \mathbb{S}^{d-1}) \lesssim \log N^- + \log N^+,$$

the result follows from Theorem 2.3. ∎

6.2 Proof of Theorem 1.4

Let $r \leq \frac{1}{C'\lambda}\delta^2$, where $C' > 0$ denotes the absolute constant from Theorem 2.3. Set $\mathcal{C}^+ := \{c_1^+, \ldots, c_{N^+}^+\}$ and $\mathcal{C}^- := \{c_1^-, \ldots, c_{N^-}^-\}$. The sets $\mathcal{C}^+, \mathcal{C}^- \subset \mathbb{B}_2^d$ are δ-separated and the sets $\mathcal{X}_j^+ := \mathcal{X}^+ \cap \mathbb{B}_2^d(c_j^+, r)$ for $j \in [N^+]$, and $\mathcal{X}_l^- := \mathcal{X}^- \cap \mathbb{B}_2^d(c_l^-, r)$ for $l \in [N^-]$, cover $\mathcal{X}^+$ and $\mathcal{X}^-$, respectively. Further, since $\mathcal{C}^+$ and $\mathcal{C}^-$ are δ-separated,

$$r \leq \frac{1}{C'\lambda}\operatorname{dist}^2(c_j^+, \mathcal{C}^-), \quad r \leq \frac{1}{C'\lambda}\operatorname{dist}^2(c_l^-, \mathcal{C}^+)$$

for all $j \in [N^+], l \in [N^-]$. This shows that $\mathcal{C}^+$ and $\mathcal{C}^-$ form a $C'\lambda$-mutual covering for $\mathcal{X}^+$ and $\mathcal{X}^-$. Therefore, $\mathcal{X}^+$ and $\mathcal{X}^-$ have $(1, \delta, C'\lambda)$-mutual complexity (N^+, N^-, w^+, w^-) with $w^+ = \max_{j \in [N^+]} w(\mathcal{X}_j^+)$ and $w^- = \max_{l \in [N^-]} w(\mathcal{X}_l^-)$. By Lemma 6.1,

$$w(\mathcal{X}^+) = w\Big(\bigcup_{j \in [N^+]} \mathbb{B}_2^d(c_j^+, r)\Big) \lesssim r\sqrt{d} + \sqrt{\log N^+}$$

and for any $l \in [N^-]$,

$$w(\mathcal{X}_l^-) = w(\mathcal{X}^- \cap \mathbb{B}_2^d(c_l^-, r)) \lesssim r\sqrt{d},$$

which implies $w^- \lesssim r\sqrt{d}$. Analogously, $w^+ \lesssim r\sqrt{d}$. The result now follows from Theorem 2.3 by additionally using that

$$w^2(\operatorname{cone}(\mathcal{X}^- - \mathcal{X}^-) \cap \mathbb{S}^{d-1}) + w^2(\operatorname{cone}(\mathcal{X}^+ - \mathcal{X}^+) \cap \mathbb{S}^{d-1}) \lesssim d.$$

$\blacksquare$

6.3 Proof of Theorem 1.5

For $c > 0$ an absolute constant that is specified later, let $\mathcal{C}^+ = \{c_1^+, \ldots, c_{N^+}^+\} \subset \mathcal{X}^+$ and $\mathcal{C}^- = \{c_1^-, \ldots, c_{N^-}^-\} \subset \mathcal{X}^-$ be minimal $c\delta^2/\lambda$-coverings of $\mathcal{X}^+$ and $\mathcal{X}^-$, respectively. Then $N^+ = \mathcal{N}(\mathcal{X}^+, c\delta^2/\lambda)$ and $N^- = \mathcal{N}(\mathcal{X}^-, c\delta^2/\lambda)$. Since $\mathcal{X}^+$ and $\mathcal{X}^-$ are δ-separated, it follows that $\mathcal{C}^+$ and $\mathcal{C}^-$ are δ-separated, too. By definition of covering, the sets $\mathcal{X}_j^+ := \mathcal{X}^+ \cap \mathbb{B}_2^d(c_j^+, c\delta^2/\lambda)$ for $j \in [N^+]$, and $\mathcal{X}_l^- := \mathcal{X}^- \cap \mathbb{B}_2^d(c_l^-, c\delta^2/\lambda)$ for $l \in [N^-]$, cover $\mathcal{X}^+$ and $\mathcal{X}^-$, respectively. Further, since $\mathcal{C}^+$ and $\mathcal{C}^-$ are δ-separated,

$$c\delta^2/\lambda \leq c\lambda^{-1}\operatorname{dist}^2(c_j^+, \mathcal{C}^-), \quad c\delta^2/\lambda \leq c\lambda^{-1}\operatorname{dist}^2(c_l^-, \mathcal{C}^+)$$

for all $j \in [N^+], l \in [N^-]$. This shows that $\mathcal{C}^+$ and $\mathcal{C}^-$ form a $\frac{\lambda}{c}$-mutual covering for $\mathcal{X}^+$ and $\mathcal{X}^-$. Therefore, $\mathcal{X}^+$ and $\mathcal{X}^-$ have $(1, \delta, \frac{\lambda}{c})$-mutual complexity (N^+, N^-, w^+, w^-) with $N^+ = \mathcal{N}(\mathcal{X}^+, c\delta^2/\lambda), N^- = \mathcal{N}(\mathcal{X}^-, c\delta^2/\lambda), w^+ = w(\mathcal{X}^+)$ and $w^- = w(\mathcal{X}^-)$. Choosing $c = \frac{1}{C'}$ with C' the absolute constant from Theorem 2.3 and applying Theorem 2.3 yields the result. $\blacksquare$

Acknowledgments

References

[ABB15]   S. An, F. Boussaid, and M. Bennamoun. "How Can Deep Rectifier Networks Achieve Linear Separability and Preserve Distances?" *Proceedings of the 32nd International Conference on Machine Learning*. Ed. by F. Bach and D. Blei. Vol. 37. 2015, 514–523.

[APVZ14]   A. Andoni, R. Panigrahy, G. Valiant, and L. Zhang. "Learning polynomials with neural networks". *Proceedings of the 31st International Conference on International Conference on Machine Learning (ICML)*. 2014, 1908–1916.

[AB19]   D. Arpit and Y. Bengio. "The benefits of over-parameterization at initialization in deep ReLU networks". Preprint arXiv:1901.03611. 2019.

[Bac17]   F. Bach. "On the equivalence between kernel quadrature rules and random feature expansions". *J. Mach. Learn. Res.* 18.1 (2017), 714–751.

[BMR17]   A. S. Bandeira, D. G. Mixon, and B. Recht. "Compressive Classification and the Rare Eclipse Problem". *Compressed Sensing and its Applications: Second International MATHEON Conference 2015*. Ed. by H. Boche, G. Caire, R. Calderbank, M. März, G. Kutyniok, and R. Mathar. Applied and Numerical Harmonic Analysis. Springer Cham, 2017, 197–220.

[BLM13]   S. Boucheron, G. Lugosi, and P. Massart. *Concentration Inequalities: A Nonasymptotic Theory of Independence*. Oxford University Press, 2013.

[BN20]   G. Bresler and D. Nagaraj. "A Corrective View of Neural Networks: Representation, Memorization and Learning". *Proceedings of Thirty Third Conference on Learning Theory (COLT)*. 2020, 848–901.

[CXJ17]   V. Cambareri, C. Xu, and L. Jacques. "The rare eclipse problem on tiles: Quantised embeddings of disjoint convex sets". *Proceedings of the 2017 International Conference on Sampling Theory and Applications (SampTA)*. 2017.

[CRPW12]   V. Chandrasekaran, B. Recht, P. A. Parrilo, and A. S. Willsky. "The Convex Geometry of Linear Inverse Problems". *Found. Comput. Math.* 12.6 (2012), 805–849.

[Dir19]   S. Dirksen. "Quantized Compressed Sensing: A Survey". *Compressed Sensing and Its Applications: Third International MATHEON Conference 2017*. Ed. by H. Boche, G. Caire, R. Calderbank, G. Kutyniok, R. Mathar, and P. Petersen. Applied and Numerical Harmonic Analysis. Birkhäuser Cham, 2019, 67–95.

[DM18a]   S. Dirksen and S. Mendelson. "Non-Gaussian Hyperplane Tessellations and Robust One-Bit Compressed Sensing". *J. Eur. Math. Soc.* (2018), to appear. Preprint arXiv:1805.09409.

[DM18b]   S. Dirksen and S. Mendelson. "Robust one-bit compressed sensing with partial circulant matrices". Preprint arXiv:1812.06719. 2018.

[DS20]   S. Dirksen and A. Stollenwerk. "Binarized Johnson-Lindenstrauss embeddings". Preprint arXiv:2009.08320. 2020.

[FR13]   S. Foucart and H. Rauhut. *A Mathematical Introduction to Compressive Sensing*. Applied and Numerical Harmonic Analysis. Birkhäuser Basel, 2013.

[GM04]   A. A. Giannopoulos and V. D. Milman. "Asymptotic Convex Geometry Short Overview". *Different Faces of Geometry*. Ed. by S. Donaldson, Y. Eliashberg, and M. Gromov. Springer Boston, 2004, 87–162.

[GSB16]   R. Giryes, G. Sapiro, and A. M. Bronstein. "Deep neural networks with random gaussian weights: A universal classification strategy?" *IEEE Trans. Signal Process.* 64.13 (2016), 3444–3457.

[GSB20]   R. Giryes, G. Sapiro, and A. M. Bronstein. "Corrections to: "Deep Neural Networks With Random Gaussian Weights: A Universal Classification Strategy?"" *IEEE Trans. Signal Process.* 68 (2020), 529–531.

[GBC16]   I. Goodfellow, Y. Bengio, and A. Courville. *Deep Learning*. MIT Press, 2016.

[Gor88]   Y. Gordon. "On Milman's inequality and random subspaces which escape through a mesh in $\mathbb{R}^n$". *Geometric Aspects of Functional Analysis*. Ed. by J. Lindenstrauss and V. D. Milman. Vol. 1317. Lecture Notes in Mathematics. Springer Berlin Heidelberg, 1988, 84–106.

[HZRS15]   K. He, X. Zhang, S. Ren, and J. Sun. "Delving deep into rectifiers: Surpassing human-level performance on imagenet classification". *Proceedings of the 2015 IEEE International Conference on Computer Vision (ICCV)*. 2015, 1026–1034.

[HSSV21]   D. Hsu, C. Sanford, R. A. Servedio, and E.-V. Vlatakis-Gkaragkounis. "On the Approximation Power of Two-Layer Networks of Random ReLUs". Preprint arXiv:2102.02336. 2021.

[HZS06]   G.-B. Huang, Q.-Y. Zhu, and C.-K. Siew. "Extreme learning machine: theory and applications". *Neurocomputing* 70.1-3 (2006), 489–501.

[JC17]   L. Jacques and V. Cambareri. "Time for dithering: fast and quantized random embeddings via the restricted isometry property". *Inf. Inference* 6.4 (2017), 441–476.

[JLBB13]   L. Jacques, J. N. Laska, P. T. Boufounos, and R. G. Baraniuk. "Robust 1-Bit Compressive Sensing via Binary Stable Embeddings of Sparse Vectors". *IEEE Trans. Inf. Theory* 59.4 (2013), 2082–2102.

[JMPS21]   H. C. Jung, J. Maly, L. Palzer, and A. Stollenwerk. "Quantized Compressed Sensing by Rectified Linear Units". *IEEE Trans. Inf. Theory* 67.6 (2021), 4125–4149.

[LBH15]   Y. LeCun, Y. Bengio, and G. Hinton. "Deep learning". *Nature* 521.7553 (2015), 436–444.

[LTOS21]   Z. Li, J.-F. Ton, D. Oglic, and D. Sejdinovic. "Towards a Unified Analysis of Random Fourier Features". *J. Mach. Learn. Res.* 22.108 (2021), 1–51.

[LHCS20]   F. Liu, X. Huang, Y. Chen, and J. A. K. Suykens. "Random Features for Kernel Approximation: A Survey on Algorithms, Theory, and Beyond". Preprint arXiv:2004.11154. 2020.

[Men16]   S. Mendelson. "Upper bounds on product and multiplier empirical processes". *Stoch. Proc. Appl.* 126.12 (2016), 3652–3680.

[NNSS20]   D. Needell, A. A. Nelson, R. Saab, and P. Salanevich. "Random Vector Functional Link Networks for Function Approximation on Manifolds". Preprint arXiv:2007.15776. 2020.

[OR15]   S. Oymak and B. Recht. "Near-Optimal Bounds for Binary Embeddings of Arbitrary Sets". Preprint arXiv:1512.04433. 2015.

[PV14]   Y. Plan and R. Vershynin. "Dimension Reduction by Random Hyperplane Tessellations". *Discrete Comput. Geom.* 51.2 (2014), 438–461.

[RR07]   A. Rahimi and B. Recht. "Random features for large-scale kernel machines". *Advances in Neural Information Processing Systems 20*. 2007, 1177–1184.

[RR08]   A. Rahimi and B. Recht. "Weighted Sums of Random Kitchen Sinks: Replacing minimization with randomization in learning". *Advances in Neural Information Processing Systems 21*. 2008, 1313–1320.

[RR17]   A. Rudi and L. Rosasco. "Generalization Properties of Learning with Random Features". *Proceedings of the 31st International Conference on Neural Information Processing Systems*. NIPS'17. Long Beach, California, USA: Curran Associates Inc., 2017, 3218–3228. ISBN: 9781510860964.

[Sax+11]   A. M. Saxe, P. W. Koh, Z. Chen, M. Bhand, B. Suresh, and A. Y. Ng. "On random weights and unsupervised feature learning". *Proceedings of the 28th International Conference on International Conference on Machine Learning (ICML)*. 2011, 1089–1096.

[Sch15]   J. Schmidhuber. "Deep learning in neural networks: An overview". *Neural Netw.* 61 (2015), 85–117.

[SB14]   S. Shalev-Shwartz and S. Ben-David. *Understanding machine learning: From theory to algorithms*. Cambridge University Press, 2014.

[SC08]   I. Steinwart and A. Christmann. *Support Vector Machines*. Information Science and Statistics. Springer, 2008.

[SGT18]   Y. Sun, A. Gilbert, and A. Tewari. "On the Approximation Properties of Random ReLU Features". Preprint arXiv:1810.04374. 2018.

[Tal14]   M. Talagrand. *Upper and Lower Bounds for Stochastic Processes*. Vol. 3. Ergebnisse der Mathematik und ihrer Grenzgebiete. Springer Berlin Heidelberg, 2014.

[Ver18]   R. Vershynin. *High-Dimensional Probability: An Introduction with Applications in Data Science*. Vol. 47. Cambridge Series in Statistical and Probabilistic Mathematics. Cambridge University Press, 2018.

[Ver20]   R. Vershynin. "Memory Capacity of Neural Networks with Threshold and Rectified Linear Unit Activations". *SIAM J. Math. Data Sci.* 2.4 (2020), 1004–1033.

[XJ20]   C. Xu and L. Jacques. "Quantized compressive sensing with RIP matrices: the benefit of dithering". *Inf. Inference* 9.3 (2020), 543–586.

[YS19]   G. Yehudai and O. Shamir. "On the Power and Limitations of Random Features for Understanding Neural Networks". *Advances in Neural Information Processing Systems 32*. 2019.

[YSJ19]   C. Yun, S. Sra, and A. Jadbabaie. "Small ReLU networks are powerful memorizers: a tight analysis of memorization capacity". *Advances in Neural Information Processing Systems 32*. 2019.

[ZBHRV17]   C. Zhang, S. Bengio, M. Hardt, B. Recht, and O. Vinyals. "Understanding deep learning requires rethinking generalization". *International Conference on Learning Representations (ICLR)*. 2017.

[ZBHRV21]   C. Zhang, S. Bengio, M. Hardt, B. Recht, and O. Vinyals. "Understanding deep learning (still) requires rethinking generalization". *Commun. ACM* 64.3 (2021), 107–115.