

Topic 5: Uniform convergence

5.1 Uniform convergence of training error rates

5.1.1 Finite hypothesis classes

Our approach to analyzing ERM in Theorem 4.4 was to establish uniform convergence of training error rates to population error rates for hypotheses from $\mathcal{H}$: for $\mathbf{S} \sim P^T$ for any probability distribution P over $\mathcal{X} \times \mathcal{Y}$,

$$\Pr\left(\forall h \in \mathcal{H}, |\text{err}(h; \mathbf{S}) - \text{err}(h; P)| \leq \sqrt{\frac{\ln(2|\mathcal{H}|/\delta)}{T}}\right) \geq 1 - \delta. \quad (5.1)$$

In the $1 - \delta$ probability event from (5.1), a bound is established on the difference between training and population error rates that holds uniformly for all $h \in \mathcal{H}$. This should be contrasted with a statement of the following form:

$$\forall h \in \mathcal{H}, \Pr\left(|\text{err}(h; \mathbf{S}) - \text{err}(h; P)| \leq \sqrt{\frac{\ln(2/\delta)}{T}}\right) \geq 1 - \delta. \quad (5.2)$$

In each of the $1 - \delta$ probability events from (5.2), the difference between training and population error rates is bounded only for a single hypothesis. Going from (5.2) to (5.1) was just a matter of applying the union bound, replacing δ in (5.2) with $\delta/|\mathcal{H}|$. However, this is only possible if $\mathcal{H}$ is finite; for infinite $\mathcal{H}$, the simple union bound does not work.

5.1.2 Impossibility for some infinite hypothesis classes

For some infinite hypothesis classes, there is no hope to establish a uniform bound on $|\text{err}(h; \mathbf{S}) - \text{err}(h; P)|$ that tends to zero with the number of training data. For example, suppose $\mathcal{X} = [0, 1]$, $\mathcal{Y} = \{0, 1\}$, and $\mathcal{H} = \{h_V \mid V \subseteq \mathcal{X}, |V| < \infty\}$, where $h_V(x) := \mathbf{1}\{x \in V\}$ for $x \in \mathcal{X}$. And suppose $(\mathbf{x}, \mathbf{y}) \sim P$ has $\mathbf{x}$ uniformly distributed on $[0, 1]$, and $\mathbf{y} = 0$ deterministically. Then for all $h_V \in \mathcal{H}$,

$$\text{err}(h_V; P) = \Pr(h_V(\mathbf{x}) = 1) = \int_0^1 \mathbf{1}\{x \in V\} dx = 0.$$

However, with $\mathbf{S} = ((\mathbf{x}_1, \mathbf{y}_1), \dots, (\mathbf{x}_T, \mathbf{y}_T)) \sim P^T$ and $\mathbf{V} := \{\mathbf{x}_1, \dots, \mathbf{x}_T\}$, we have

$$\text{err}(h_{\mathbf{V}}; \mathbf{S}) = 1.$$

Therefore, we always have

$$\exists h_V \in \mathcal{H} \text{ s.t. } |\text{err}(h_V; \mathbf{S}) - \text{err}(h_V; P)| = 1.$$

What makes the above hypothesis class unsuitable for uniform convergence? Are there infinite hypothesis classes for which a uniform convergence guarantee like (5.1) holds?

5.2 Vapnik-Chervonenkis theory

5.2.1 Empirical covers and empirical processes

Vapnik and Chervonenkis (1971) developed a theory to study algorithms for learning binary classifiers, and a key part of their theory is a technique for proving uniform convergence of training error rates for hypothesis classes. Their basic technique has two components. The first component characterizes which function classes can be effectively replaced by a finite approximation. The relevant finite approximation is called an empirical cover, and there is a combinatorial property of a function class that determines if it always admits a non-trivial empirical cover.

The second component of the technique concerns how to effectively use a “union bound” over an empirical cover to establish uniform convergence of empirical means over an iid sample to the population means. The (random) probability distribution $\mathbf{P}_n$ that puts $1/n$ mass on each point in a size- n iid sample from P is called an empirical distribution (or empirical measure). (The average value of a function under $\mathbf{P}_n$ is called an empirical average.) The study of how $\mathbf{P}_n$ deviates from P is the subject of empirical process theory, and Vapnik and Chervonenkis drawn up techniques from that literature.

5.2.2 VC dimension

Let $\mathcal{F} \subseteq \{0, 1\}^{\mathcal{Z}}$ be a class of $\{0, 1\}$ -valued functions defined on a domain $\mathcal{Z}$. The behaviors of $\mathcal{F}$ on a finite set of m points $\{z_1, \dots, z_m\} \subseteq \mathcal{Z}$ is

$$\mathcal{F}[\{z_1, \dots, z_m\}] := \{(f(z_1), \dots, f(z_m)) \mid f \in \mathcal{F}\} \subseteq \{0, 1\}^m.$$

Intuitively, we only care about the behaviors of $\mathcal{F}$ on a finite set of points when it comes to understanding the empirical averages of functions from $\mathcal{F}$. However, the number of behaviors, even on a finite set, can vary a lot depending on the function class. What about a function class determines whether the number of behaviors is small or large?

We say that $\mathcal{F}$ shatters the points $\{z_1, \dots, z_m\}$ if

$$\mathcal{F}[\{z_1, \dots, z_m\}] = \{0, 1\}^m;$$

that is, for each $(b_1, \dots, b_m) \in \{0, 1\}^m$, there exists some $f \in \mathcal{F}$ such that $f(z_i) = b_i$ for all $i \in [m]$. The Vapnik-Chervonenkis (VC) dimension of $\mathcal{F}$, denoted $\text{VCdim}(\mathcal{F})$, is the largest value m such that there exists a set of m points from $\mathcal{Z}$ that is shattered by $\mathcal{F}$; if there are arbitrarily large sets of point shattered by $\mathcal{F}$, then $\text{VCdim}(\mathcal{F}) = \infty$.

Lemma 5.1. For any $\mathcal{F} \subseteq \{0, 1\}^{\mathcal{Z}}$, $\text{VCdim}(\mathcal{F}) \leq \log_2(|\mathcal{F}|)$.

Proof. If a set $\{z_1, \dots, z_m\}$ of cardinality m is shattered by $\mathcal{F}$, then $|\mathcal{F}[\{z_1, \dots, z_m\}]| = 2^m$. Now map every $f \in \mathcal{F}$ to $(f(z_1), \dots, f(z_m)) \in \mathcal{F}[\{z_1, \dots, z_m\}]$ to deduce $|\mathcal{F}| \geq 2^m$ as well. $\square$

Example 5.1. Recall the class of monotone disjunctions MDISJ_n over the Boolean n -cube $\{0, 1\}^n$:

$$\text{MDISJ}_n := \{d_S \mid S \subseteq [n]\}$$

where $d_S(x_1, \dots, x_n) = \bigvee_{j \in S} x_j$. We claim that $\text{VCdim}(\text{MDISJ}_n) = n$. We have $\text{VCdim}(\text{MDISJ}_n) \leq n$ by Lemma 5.1 since $|\text{MDISJ}_n| = 2^n$. So it remains to show that $\text{VCdim}(\text{MDISJ}_n) \geq n$.

We claim that the set $T := \{e_1, \dots, e_n\}$ is shattered by MDISJ_n ; here, $e_i = (e_i[1], \dots, e_i[n]) \in \{0, 1\}^n$ is the vector that is all-0's except in the i -th position (i.e., $e_i[j] = \mathbb{1}\{i = j\}$). To show this, we must establish that every possible labeling of T can be realized by functions from MDISJ_n . Consider an arbitrary labeling $b \in \{0, 1\}^n$, written $b = (b_1, \dots, b_n)$; b_i is the desired label of e_i . Set $S := \{j \in [n] \mid b_j = 1\}$. Then for each $i \in [n]$,

$$d_S(e_i) = \bigvee_{j \in S} e_i[j] = \bigvee_{j \in S} \mathbb{1}\{i = j\} = \mathbb{1}\{i \in S\} = b_i.$$

Therefore the labeling b is realized by the monotone disjunction d_S . We conclude that T is shattered by MDISJ_n , and hence $\text{VCdim}(\text{MDISJ}_n) \geq |T| = n$. $\diamond$

Example 5.2. Recall the class of linear threshold functions over Euclidean n -space $\mathbb{R}^n$:

$$\text{LTF}_{\mathbb{R}^n} := \left\{ f: \mathbb{R}^n \rightarrow \{0, 1\} \mid \exists (w, \theta) \in \mathbb{R}^n \times \mathbb{R} \text{ s.t. } f(x) = \mathbb{1}\{\langle w, x \rangle > \theta\} \right\}.$$

We claim that $\text{VCdim}(\text{LTF}_{\mathbb{R}^n}) = n + 1$.

We first show $\text{VCdim}(\text{LTF}_{\mathbb{R}^n}) \geq n + 1$ by exhibiting a set of $n + 1$ points that is shattered by $\text{LTF}_{\mathbb{R}^n}$. We claim that the set $T := \{0, e_1, \dots, e_n\}$, where e_i is the i -th elementary basis vector in $\mathbb{R}^n$, is shattered by $\text{LTF}_{\mathbb{R}^n}$. To show this, we must establish that every possible labeling of T can be realized by functions from $\text{LTF}_{\mathbb{R}^n}$. Consider an arbitrary labeling $b \in \{0, 1\}^{n+1}$, written $b = (b_0, b_1, \dots, b_n)$; b_0 is the desired label of 0, and for each $i \in [n]$, b_i is the desired label of e_i . Let $\tilde{b} := 2b - 1 = (\tilde{b}_0, \tilde{b}_1, \dots, \tilde{b}_n) \in \{-1, 1\}^{n+1}$ be the $\{-1, 1\}$ -version of b , and consider the weight vector w and threshold θ given by

$$w := 2\tilde{b}_1 e_1 + \dots + 2\tilde{b}_n e_n, \quad \theta := -\tilde{b}_0$$

Then

$$\mathbb{1}\{\langle w, 0 \rangle > \theta\} = \mathbb{1}\{0 > -\tilde{b}_0\} = b_0,$$

and for each $i \in [n]$,

$$\mathbb{1}\{\langle w, e_i \rangle > \theta\} = \mathbb{1}\{2\tilde{b}_i > -\tilde{b}_0\} = \mathbb{1}\{\tilde{b}_i + \tilde{b}_0/2 > 0\} = b_i.$$

Therefore the labeling b is realized by the LTF $x \mapsto \mathbb{1}\{\langle w, x \rangle > \theta\}$. We conclude that T is shattered by $\text{LTF}_{\mathbb{R}^n}$, and hence $\text{VCdim}(\text{LTF}_{\mathbb{R}^n}) \geq |T| = n + 1$.

It now remains to show that $\text{VCdim}(\text{LTF}_{\mathbb{R}^n}) < n + 2$. To do this, we use the following simple observation (recalling the discussion in Section 3.3.3 about Perceptron): if a set $T \subset \mathbb{R}^n$ is shattered by $\text{LTF}_{\mathbb{R}^n}$, then $\tilde{T} := \{(x, -1) \mid x \in T\} \subset \mathbb{R}^{n+1}$ is shattered by $\text{HLTF}_{\mathbb{R}^{n+1}}$. So it suffices to show that no set of $n + 2$ points is shattered by $\text{HLTF}_{\mathbb{R}^{n+1}}$. Consider any $n + 2$ points $x_1, \dots, x_{n+2} \in \mathbb{R}^{n+1}$. They must be linearly dependent, so we can write

$$c_1 x_1 + \dots + c_{n+2} x_{n+2} = 0$$

for some $c_1, \dots, c_{n+2} \in \mathbb{R}$ that are not all zero. In particular,

$$\langle w, c_1 x_1 + \dots + c_{n+2} x_{n+2} \rangle = 0, \quad \forall w \in \mathbb{R}^{n+1}. \quad (5.3)$$

We may assume without loss of generality that $c_i > 0$ for some $i \in [n + 2]$ (or else replace all c_i with $-c_i$). Define $b = (b_1, \dots, b_{n+2}) \in \{0, 1\}^{n+2}$ by $b_i := \mathbb{1}\{c_i > 0\}$ for each $i \in [n + 2]$. Suppose, for sake

of contradiction, that the labeling b is realized by some HLTF, i.e., there exists $w \in \mathbb{R}^{n+1}$ such that, for all $i \in [n+2]$,

$$\langle w, x_i \rangle > 0 \iff c_i > 0.$$

This implies

$$\left\langle w, \sum_{i=1}^{n+2} c_i x_i \right\rangle = \sum_{i=1}^{n+2} c_i \langle w, x_i \rangle > 0,$$

which is a contradiction of (5.3). So the labeling b of $x_1, \dots, x_{n+2}$ cannot be realized by any HLTF. Since this holds for any choice of $n+2$ points from $\mathbb{R}^{n+1}$, it follows that no set of $n+2$ points from $\mathbb{R}^{n+1}$ is shattered by $\text{HLTF}_{\mathbb{R}^{n+1}}$. This means that no set of $n+2$ points from $\mathbb{R}^n$ is shattered by $\text{LTF}_{\mathbb{R}^n}$, and hence $\text{VCdim}(\text{LTF}_{\mathbb{R}^n}) < n+2$. $\diamond$

Exercise 5.1. Let $\mathcal{X} := \{0, 1, 2, \dots, N-1\}$ for $N \in \mathbb{N}$, and

$$\mathcal{H} := \{h: \mathcal{X} \rightarrow \{0, 1\} \mid \exists \theta \in \mathbb{R} \text{ s.t. } h(x) = \mathbb{1}\{x > \theta\} \forall x \in \mathcal{X}\}.$$

What are $\text{Ldim}(\mathcal{H})$ and $\text{VCdim}(\mathcal{H})$? $\diamond$

Exercise 5.2. Let $\mathcal{H} \subseteq \{0, 1\}^{\mathcal{X}}$ denote a hypothesis class on the domain $\mathcal{X}$. Let $\mathcal{H}^* := \{e_x \mid x \in \mathcal{X}\}$ be the dual class where

$$e_x(h) := h(x) \quad \forall h \in \mathcal{H}.$$

(a) Show that

$$\text{VCdim}(\mathcal{H}^*) \leq 2^{\text{VCdim}(\mathcal{H})+1} - 1.$$

(b) Let $\mathcal{X} := \mathbb{R}^n$ and $\mathcal{H} := \{h: \mathbb{R}^n \rightarrow \{0, 1\} \mid \exists w \in \mathbb{R}^n \text{ s.t. } h(x) = \mathbb{1}\{\langle w, x \rangle > 0\} \forall x \in \mathbb{R}^n\}$. Show that $\text{VCdim}(\mathcal{H}^*) = \text{VCdim}(\mathcal{H})$. $\diamond$

5.2.3 Sauer's lemma

If $\mathcal{F}$ has VC dimension d , then (by definition) there are finite point sets $T \subset \mathcal{Z}$ of all sizes $|T| \leq d$ for which the behaviors of $\mathcal{F}$ on T is all behaviors: $|\mathcal{F}[T]| = 2^n$. However, if $d < \infty$, this exponential growth stops as soon as $|T| \geq d$, after which the number of behaviors grows at most polynomially.

Define the growth function $\Pi_{\mathcal{F}}: \mathbb{N} \rightarrow \mathbb{N}$ for the function class $\mathcal{F}$ by

$$\Pi_{\mathcal{F}}(n) := \max\{|\mathcal{F}[T]| \mid T \subseteq \mathcal{Z}, |T| = n\}, \quad \forall n \in \mathbb{N}.$$

Sauer's lemma, given in Lemma 5.2, implies that $\Pi_{\mathcal{F}}$ is asymptotically upper-bounded by a polynomial in n whenever $\text{VCdim}(\mathcal{F}) < \infty$.

Lemma 5.2 (Sauer, 1972; Shelah, 1972). *If $d := \text{VCdim}(\mathcal{F}) < \infty$, then $\Pi_{\mathcal{F}}(n) \leq \binom{n}{\leq d} := \sum_{k=0}^d \binom{n}{k}$.*

There are a number of proofs of Lemma 5.2. The proof we give below is based on those of Alon (1983) and Frankl (1983) using a technique called shifting.

Lemma 5.2 concerns subsets of the Boolean n -cube $\{0, 1\}^n$, so it is useful to specialize the notions of shattering and VC dimension to such subsets. For a subset $V \subseteq \{0, 1\}^n$ and $J \subseteq [n]$, let $V|_J := \{(v_i)_{i \in J} \mid v \in V\}$. We say V shatters J if $V|_J = \{0, 1\}^J$, and define $\text{VCdim}(V)$ to be the largest value d such that there exists $J \subseteq [n]$ of cardinality $|J| = d$ such that V shatters J .

Note that each $v \in \{0, 1\}^n$ can also be regarded as the characteristic vector for a subset of $[n]$; we will write v to also refer to that subset. With this overloaded notation, we have that $V \subseteq 2^{[n]}$ shatters $J \subseteq [n]$ if and only if $\{v \cap J \mid v \in V\} = 2^J$. When v is regarded this way, we also use the shorthand $v - i$ to mean $v \setminus \{i\}$ for any $i \in [n]$.




Figure 5.1: Shifting $V \subseteq \{0, 1\}^n$ (for $n = 4$) along coordinate $i = 1$. The family V is indicated by the gray colored vertices. (Vertices corresponding to $v \in \{0, 1\}^n \setminus V$ are also depicted but colored white.) The dark solid arrows indicate $v \in V$ that are successfully shifted (e.g., $S_{1,V}((1, 1, 0, 0)) = (0, 1, 0, 0)$). Successful shifts move $v \in V$ down to lower levels of the Boolean n -cube. The dark dashed arrow indicates an unsuccessful attempt to shift $v = (1, 0, 0, 0)$ that is blocked by the presence of $(0, 0, 0, 0) \in V$.

Proposition 5.1. *For any finite set $T \subseteq \mathcal{Z}$, we have $\text{VCdim}(\mathcal{F}[T]) \leq \text{VCdim}(\mathcal{F})$.*

Exercise 5.3. Prove Proposition 5.1, and explain in what sense the inequality is tight. $\diamond$

The goal of the proof of Lemma 5.2 is to perform some transformations to $V := \mathcal{F}[T]$ (for some $T \subseteq \mathcal{Z}$ with $|T| = n$) that (i) preserve the cardinality of V , and (ii) do not increase the VC dimension of V . After all of the transformations, the resulting family $W \subseteq 2^{[n]}$ will have some useful properties that make it very easy to bound $|W|$ (and hence also $|V|$).

The transformations to V are called (down)shifting along a coordinate $i \in [n]$. The shifting transformations are as follows.

$$\text{Shift } v \in V \text{ along coordinate } i \in [n]: \quad S_{i,V}(v) := \begin{cases} v - i & \text{if } i \in v \text{ and } v - i \notin V, \\ v & \text{otherwise.} \end{cases}$$

$$\text{Shift } V \text{ along coordinate } i \in [n]: \quad S_i(V) := \{S_{i,V}(v) \mid v \in V\}.$$

If $S_{i,V}(v) = v - i \neq v$, then we say that v is successfully shifted along coordinate i . See Figure 5.1 for an illustration. There are two possible reasons for an unsuccessful shift along coordinate i . First, it could be that $i \notin v$. Second, it could be that $i \in v$ but $v - i \in V$; in this case, the attempted shift is “blocked” by the presence of $v - i \in V$.

Exercise 5.4. For the $V \subseteq \{0, 1\}^4$ from Figure 5.1, list the members of $S_2(V)$. $\diamond$

Proposition 5.2. *For any $V \subseteq 2^{[n]}$ and any $i \in [n]$:*

1. $|S_i(V)| = |V|$.

2. For any $J \subseteq [n]$, if J is shattered by $S_i(V)$, then J is shattered by V . (Consequently, $\text{VCdim}(S_i(V)) \leq \text{VCdim}(V)$.)

Proof. Observe that $v \in V$ is successfully shifted by $S_{i,V}$ if and only if $v - i \neq V$; so V and $S_i(V)$ are in 1-to-1 correspondence (via $S_{i,V}$). This proves the first item of the claim.

To prove the second item, consider any $J \subseteq [n]$ shattered by $S_i(V)$, so $\{S_{i,V}(v) \cap J \mid v \in V\} = 2^J$. We need to show that J is also shattered by V . There are two cases to consider.

- Case 1: $i \notin J$. Then for every $v \in V$, we have $S_{i,V}(v) \cap J = v \cap J$ since $S_{i,V}$ only (possibly) removes i from v . So $\{v \cap J \mid v \in V\} = \{S_{i,V}(v) \cap J \mid v \in V\} = 2^J$, which shows that J is also shattered by V .
- Case 2: $i \in J$. Since $\{S_{i,V}(v) \cap J \mid v \in V\} = 2^J$, for any choice of $K \in 2^J$, there exists $v \in V$ such that

$$S_{i,V}(v) \cap J = K \cup \{i\}.$$

This means that $i \in S_{i,V}(v)$, implying that $S_{i,V}(v) = v$, an unsuccessful shift of $v \in V$ along coordinate i , blocked by the presence of $u := v - i$ in V . Therefore, we have both

$$\begin{aligned} u \cap J &= K \setminus \{i\} \in V \quad \text{and} \\ v \cap J &= K \cup \{i\} \in V. \end{aligned}$$

Since this holds for all choices of $K \subseteq J$, it follows that V also shatters J . □

Exercise 5.5. Can there exist $V \subseteq \{0,1\}^n$ such that $\text{VCdim}(S_1(V)) < \text{VCdim}(V)$? ◇

Suppose, starting with V , we repeatedly apply S_i to V for each $i \in [n]$. Whenever there is a successful shift, the “total weight” of V (i.e., $\sum_{v \in V} |v|$) decreases by at least 1. So this process must eventually stop with some set $W \subseteq 2^{[n]}$ with $|W| = |V|$ and $\text{VCdim}(W) \leq \text{VCdim}(V)$.

Proposition 5.3. Suppose $W \subseteq 2^{[n]}$ satisfies $S_i(W) = W$ for each $i \in [n]$. For each $v \in W$, we have $u \in W$ for all $u \subseteq v$ (i.e., W is downward-closed).

Proof. Pick any nonempty $v \in W$ and any $i \in v$. If $v - i \notin W$, then $S_{i,W}(v)$ would be a successful shift, a contradiction of the assumption that $S_i(W) = W$. Therefore it must be that $v - i \in W$. Repeatedly applying this argument shows that every subset $u \subseteq v$ is also in W . □

Proposition 5.4. Suppose $W \subseteq 2^{[n]}$ is downward-closed, and let $d := \text{VCdim}(W)$. Then

$$|W| \leq \binom{n}{\leq d}.$$

Proof. We first claim that every $v \in W$ has $|v| \leq d$. To see this, consider any $v \in W$, and let $J := v$. Since W is downward-closed, W shatters J . Since $\text{VCdim}(W) = d$, it must be that $|J| \leq d$. This establishes $|v| \leq d$.

So, every $v \in W$ is a subset of $[n]$ of cardinality at most d ; and there are at most $\binom{n}{\leq d}$ subsets of $[n]$ of cardinality at most d . □

Combining Propositions 5.1 to 5.4 (along with the above observation that the number of successful shifts is finite) completes the proof Lemma 5.2.

5.2.4 Population and empirical averages

Let P be a probability distribution over $\mathcal{Z}$ (which we regard as the “population”), and let $\mathbf{z}_{1:n} := (\mathbf{z}_1, \dots, \mathbf{z}_n) \sim P^n$ be a size- n iid sample from P . For a real-valued function $f: \mathcal{Z} \rightarrow \mathbb{R}$, we are concerned with two different “averages”:

- the population average of f

$$\mathbb{E}[f(\mathbf{z})]$$

(where $\mathbf{z} \sim P$);

- the empirical average of f

$$\frac{1}{n} \sum_{i=1}^n f(\mathbf{z}_i).$$

Our goal is to show the following theorem.

Theorem 5.1. *Let $\mathcal{F} \subseteq \{0, 1\}^{\mathcal{Z}}$ be a function class with $d := \text{VCdim}(\mathcal{F}) < \infty$. Let P be a probability distribution over $\mathcal{Z}$ and $(\mathbf{z}_1, \dots, \mathbf{z}_n) \sim P^n$. Then*

$$\mathbb{E} \left[\sup_{f \in \mathcal{F}} \left| \frac{1}{n} \sum_{i=1}^n f(\mathbf{z}_i) - \mathbb{E}[f(\mathbf{z})] \right| \right] \leq 2 \sqrt{\frac{2 \ln(2 \binom{n}{\leq d})}{n}}.$$

A proof that has a lot of overlap with that of Theorem 5.1 establishes the following theorem.

Theorem 5.2 (Vapnik and Chervonenkis, 1971). *Let $\mathcal{F} \subseteq \{0, 1\}^{\mathcal{Z}}$ be a function class with $d := \text{VCdim}(\mathcal{F}) < \infty$. Let P be a probability distribution over $\mathcal{Z}$ and $(\mathbf{z}_1, \dots, \mathbf{z}_n) \sim P^n$. Then for any $\delta \in (0, 1)$, with probability at least $1 - \delta$,*

$$\sup_{f \in \mathcal{F}} \left| \frac{1}{n} \sum_{i=1}^n f(\mathbf{z}_i) - \mathbb{E}[f(\mathbf{z})] \right| \leq O \left(\sqrt{\frac{\log(\binom{n}{\leq d} / \delta)}{n}} \right).$$

It is also possible to derive Theorem 5.2 as an immediate corollary of Theorem 5.1 by applying a probability inequality called McDiarmid’s inequality. We focus on Theorem 5.1 because its proof is more streamlined than the direct proof of Theorem 5.2 but retains most of the key ideas.

In Theorems 5.1 and 5.2, because $\mathcal{F}$ has finite VC dimension, we might hope that we can replace the supremum over $\mathcal{F}$ with a maximum over the set of “behaviors” of $\mathcal{F}$ on $\{\mathbf{z}_1, \dots, \mathbf{z}_n\}$. The thought would be that (at least for Theorem 5.2) because there are only $\binom{n}{\leq d}$ behaviors, we can establish a probability inequality for each behavior separately, and then apply a union bound. This is almost how the proofs of Theorems 5.1 and 5.2 work, but some important issues arise that require careful attention.

The first issue is that, even though $\mathcal{F}$ has only finitely-many behaviors on the size- n iid sample—and hence finitely-many values of the empirical average $\frac{1}{n} \sum_{i=1}^n f(\mathbf{z}_i)$ as f ranges over $\mathcal{F}$ —there may be infinitely-many values of the population average $\mathbb{E}[f(\mathbf{z})]$ as f ranges over $\mathcal{F}$.

A way to resolve this issue is based on the following observation: if $\mathbf{x}$ and $\mathbf{x}'$ are iid real-valued random variables, then $\mathbb{E}[\mathbf{x}] = \mathbb{E}[\mathbf{x}' | \mathbf{x}]$, which implies $\mathbb{E}[|\mathbf{x} - \mathbb{E}[\mathbf{x}]|] \leq \mathbb{E}[|\mathbf{x} - \mathbf{x}'|]$.¹ This suggests

¹In fact, if we considered squared differences instead of absolute differences, the left-hand side would be equal to half the right-hand side: $\mathbb{E}[(\mathbf{x} - \mathbb{E}[\mathbf{x}])^2] = \frac{1}{2} \mathbb{E}[(\mathbf{x} - \mathbf{x}')^2]$.

that, for the purposes of analysis only, it may be useful to consider a second size- n iid sample $\mathbf{z}'_{1:n} := (\mathbf{z}'_1, \dots, \mathbf{z}'_n) \sim P^n$ that is independent of—but identically distributed as—the original sample $\mathbf{z}_{1:n}$. We call this second sample the *ghost sample* $\mathfrak{G}$. If we could replace $\mathbb{E}[f(\mathbf{z})]$ with the empirical average of f on the ghost sample, then we could again hope to use the fact that the (ghost) empirical average takes only finitely-many values as f ranges over $\mathcal{F}$. (It is still not clear how to exactly take advantage of this, but intuitively this could seem like progress.)

Proposition 5.5 (Symmetrization by ghost sample). *Let $\mathbf{z}_{1:n}$ and $\mathbf{z}'_{1:n}$ be independent size- n iid samples from a probability distribution P over $\mathcal{Z}$. For any $\mathcal{F} \subseteq \{0, 1\}^{\mathcal{Z}}$,*

$$\mathbb{E} \left[\sup_{f \in \mathcal{F}} \left| \frac{1}{n} \sum_{i=1}^n f(\mathbf{z}_i) - \mathbb{E}[f(\mathbf{z})] \right| \right] \leq \mathbb{E} \left[\sup_{f \in \mathcal{F}} \left| \frac{1}{n} \sum_{i=1}^n f(\mathbf{z}_i) - \frac{1}{n} \sum_{i=1}^n f(\mathbf{z}'_i) \right| \right].$$

Proof. Define

$$\Delta f := \frac{1}{n} \sum_{i=1}^n f(\mathbf{z}_i) - \frac{1}{n} \sum_{i=1}^n f(\mathbf{z}'_i)$$

to be the difference between the two empirical averages of $f: \mathcal{Z} \rightarrow \mathbb{R}$. Consider any $f \in \mathcal{F}$, and observe that

$$\mathbb{E}[f(\mathbf{z})] = \mathbb{E} \left[\frac{1}{n} \sum_{i=1}^n f(\mathbf{z}_i) \right] = \mathbb{E} \left[\frac{1}{n} \sum_{i=1}^n f(\mathbf{z}'_i) \mid \mathbf{z}_{1:n} \right]$$

since $\mathbf{z}_{1:n}$ and $\mathbf{z}'_{1:n}$ are independent and identically distributed. Therefore,

$$\left| \frac{1}{n} \sum_{i=1}^n f(\mathbf{z}_i) - \mathbb{E}[f(\mathbf{z})] \right| = |\mathbb{E}[\Delta f \mid \mathbf{z}_{1:n}]| \leq \mathbb{E}[|\Delta f| \mid \mathbf{z}_{1:n}]$$

where the inequality follows by Jensen's inequality ($\mathbb{E}[\mathbf{x}]^2 \leq \mathbb{E}[\mathbf{x}^2]$). We also have

$$|\Delta f| \leq \sup_{f' \in \mathcal{F}} |\Delta f'|,$$

and therefore

$$\mathbb{E}[|\Delta f| \mid \mathbf{z}_{1:n}] \leq \mathbb{E} \left[\sup_{f' \in \mathcal{F}} |\Delta f'| \mid \mathbf{z}_{1:n} \right].$$

Since $f \in \mathcal{F}$ was arbitrary, we have

$$\sup_{f \in \mathcal{F}} \left| \frac{1}{n} \sum_{i=1}^n f(\mathbf{z}_i) - \mathbb{E}[f(\mathbf{z})] \right| \leq \mathbb{E} \left[\sup_{f' \in \mathcal{F}} |\Delta f'| \mid \mathbf{z}_{1:n} \right].$$

Taking expectation of both sides gives the claim. $\square$

Exercise 5.6. In the same setting as Proposition 5.5, show that

$$\mathbb{E} \left[\sup_{f \in \mathcal{F}} \left| \frac{1}{n} \sum_{i=1}^n f(\mathbf{z}_i) - \mathbb{E}[f(\mathbf{z})] \right| \right] \geq \frac{1}{2} \mathbb{E} \left[\sup_{f \in \mathcal{F}} \left| \frac{1}{n} \sum_{i=1}^n f(\mathbf{z}_i) - \frac{1}{n} \sum_{i=1}^n f(\mathbf{z}'_i) \right| \right]. \quad \diamond$$

Proposition 5.5 is called a “symmetrization” argument, because it effectively replaces each (possibly) asymmetric random variable $f(\mathbf{z}_i) - \mathbb{E}[f(\mathbf{z})]$ with a symmetric random variable $f(\mathbf{z}_i) - f(\mathbf{z}'_i)$. So, to prove Theorem 5.1, it suffices bound the right-hand side in Proposition 5.5. (In view of Exercise 5.6, at most a constant factor is lost.) But this raises the second issue: how to take advantage of the finitely-many behaviors of $\mathcal{F}$ on the (two) iid samples. Note that the set of behaviors is itself a random subset of $\{0, 1\}^{2n}$; what leverage is gotten by fixing this set of behaviors?

The expression on the right-hand side in Proposition 5.5 concerns two independent iid samples, $\mathbf{z}_{1:n}$ and $\mathbf{z}'_{1:n}$, which we can imagine are derived by starting with a single size- $2n$ iid sample, say, $(\mathbf{z}_1^a, \mathbf{z}_1^b, \dots, \mathbf{z}_n^a, \mathbf{z}_n^b) \sim P^{2n}$, and then, for each $i \in [n]$, assigning one of $\mathbf{z}_i^a$ and $\mathbf{z}_i^b$ to the original sample, and then assigning the other to the ghost sample. Any way of making this assignment is valid as long as it does not depend on the size- $2n$ iid sample. This flexibility turns out to be key.

Proposition 5.6. *Let $\mathbf{z}_{1:n}$ and $\mathbf{z}'_{1:n}$ be independent size- n iid samples from a probability distribution P over $\mathcal{Z}$. For any $\mathcal{F} \subseteq \{0, 1\}^{\mathcal{Z}}$ and any $(\sigma_1, \dots, \sigma_n) \in \{-1, 1\}^n$,*

$$\mathbb{E} \left[\sup_{f \in \mathcal{F}} \left| \frac{1}{n} \sum_{i=1}^n (f(\mathbf{z}_i) - f(\mathbf{z}'_i)) \right| \right] \leq 2 \mathbb{E} \left[\sup_{f \in \mathcal{F}} \left| \frac{1}{n} \sum_{i=1}^n \sigma_i f(\mathbf{z}_i) \right| \right].$$

Proof. Let $(\mathbf{z}_1^a, \mathbf{z}_1^b, \dots, \mathbf{z}_n^a, \mathbf{z}_n^b) \sim P^{2n}$ be a size- $2n$ iid sample from P . We determine $\mathbf{z}_{1:n}$ and $\mathbf{z}'_{1:n}$ as follows. For each $i \in [n]$:

- If $\sigma_i = +1$, then $\mathbf{z}_i := \mathbf{z}_i^a$ and $\mathbf{z}'_i := \mathbf{z}_i^b$.
- If $\sigma_i = -1$, then $\mathbf{z}_i := \mathbf{z}_i^b$ and $\mathbf{z}'_i := \mathbf{z}_i^a$.

Then $\mathbf{z}_{1:n}$ and $\mathbf{z}'_{1:n}$ are independent and each is a size- n iid sample from P . Moreover,

$$\sup_{f \in \mathcal{F}} \left| \frac{1}{n} \sum_{i=1}^n (f(\mathbf{z}_i) - f(\mathbf{z}'_i)) \right| = \sup_{f \in \mathcal{F}} \left| \frac{1}{n} \sum_{i=1}^n \sigma_i (f(\mathbf{z}_i^a) - f(\mathbf{z}_i^b)) \right|. \quad (5.4)$$

By the triangle inequality,

$$\sup_{f \in \mathcal{F}} \left| \frac{1}{n} \sum_{i=1}^n \sigma_i (f(\mathbf{z}_i^a) - f(\mathbf{z}_i^b)) \right| \leq \sup_{f \in \mathcal{F}} \left| \frac{1}{n} \sum_{i=1}^n \sigma_i f(\mathbf{z}_i^a) \right| + \sup_{f \in \mathcal{F}} \left| \frac{1}{n} \sum_{i=1}^n \sigma_i f(\mathbf{z}_i^b) \right|. \quad (5.5)$$

Therefore, combining (5.4) and (5.5) and taking expectations, we get

$$\begin{aligned} \mathbb{E} \left[\sup_{f \in \mathcal{F}} \left| \frac{1}{n} \sum_{i=1}^n (f(\mathbf{z}_i) - f(\mathbf{z}'_i)) \right| \right] &\leq \mathbb{E} \left[\sup_{f \in \mathcal{F}} \left| \frac{1}{n} \sum_{i=1}^n \sigma_i f(\mathbf{z}_i^a) \right| + \sup_{f \in \mathcal{F}} \left| \frac{1}{n} \sum_{i=1}^n \sigma_i f(\mathbf{z}_i^b) \right| \right] \\ &= 2 \mathbb{E} \left[\sup_{f \in \mathcal{F}} \left| \frac{1}{n} \sum_{i=1}^n \sigma_i f(\mathbf{z}_i) \right| \right] \end{aligned}$$

where the last step uses linearity of expectation and the fact that $(\mathbf{z}_1, \dots, \mathbf{z}_n)$, $(\mathbf{z}_1^a, \dots, \mathbf{z}_n^a)$, and $(\mathbf{z}_1^b, \dots, \mathbf{z}_n^b)$ have the same distribution. $\square$

At face value, it may seem that the only purpose of Proposition 5.6 is to bound a quantity involving two size- n iid samples (or a size- $2n$ iid sample) with a quantity involving only one

size- n iid sample. This is helpful, but not the main point. In fact, for any fixed assignment $(\sigma_1, \dots, \sigma_n) \in \{-1, 1\}^n$, Proposition 5.6 is a “de-symmetrization”, as it moves from the symmetric random variable $f(\mathbf{z}_i) - f(\mathbf{z}'_i)$ to the (possibly) asymmetric random variable $\sigma_i f(\mathbf{z}_i)$!

The key is that the bound in Proposition 5.6 holds for every choice of the assignment $(\sigma_1, \dots, \sigma_n)$, so the right-hand side (without the factor of 2) can be replaced by

$$\frac{1}{2^n} \sum_{\sigma \in \{-1, 1\}^n} \mathbb{E} \left[\sup_{f \in \mathcal{F}} \left| \frac{1}{n} \sum_{i=1}^n \sigma_i f(\mathbf{z}_i) \right| \right].$$

If we introduce a random sign vector $(\sigma_1, \dots, \sigma_n)$ that is (i) distributed uniformly over $\{-1, 1\}^n$, and (ii) regarded as independent of $\mathbf{z}_{1:n}$, then the above expression is identical to

$$\mathbb{E} \left[\sup_{f \in \mathcal{F}} \left| \frac{1}{n} \sum_{i=1}^n \sigma_i f(\mathbf{z}_i) \right| \right].$$

The $\sigma_1, \dots, \sigma_n$ are called iid Rademacher random variables. Combining Propositions 5.5 and 5.6 immediately gives the following proposition.

Proposition 5.7 (Symmetrization by random signs). *Let $\mathbf{z}_{1:n}$ a size- n iid sample from a probability distribution P over $\mathcal{Z}$. Let $(\sigma_1, \dots, \sigma_n)$ be n iid Rademacher random variables, independent of $\mathbf{z}_{1:n}$. For any $\mathcal{F} \subseteq \{0, 1\}^{\mathcal{Z}}$,*

$$\mathbb{E} \left[\sup_{f \in \mathcal{F}} \left| \frac{1}{n} \sum_{i=1}^n f(\mathbf{z}_i) - \mathbb{E}[f(\mathbf{z})] \right| \right] \leq 2 \mathbb{E} \left[\sup_{f \in \mathcal{F}} \left| \frac{1}{n} \sum_{i=1}^n \sigma_i f(\mathbf{z}_i) \right| \right].$$

The last step is to use the tower property of conditional expectations:

$$\mathbb{E} \left[\sup_{f \in \mathcal{F}} \left| \frac{1}{n} \sum_{i=1}^n \sigma_i f(\mathbf{z}_i) \right| \right] = \mathbb{E} \left[\mathbb{E} \left[\sup_{f \in \mathcal{F}} \left| \frac{1}{n} \sum_{i=1}^n \sigma_i f(\mathbf{z}_i) \right| \middle| \mathbf{z}_{1:n} \right] \right], \quad (5.6)$$

and now focus on the inner conditional expectation. In this conditional expectation, the iid sample $\mathbf{z}_{1:n}$ is fixed. We now (finally!) regard the supremum over $\mathcal{F}$ as a maximum over the behaviors $\mathcal{F}[\{\mathbf{z}_1, \dots, \mathbf{z}_n\}] \subseteq \{0, 1\}^n$, and we apply the following lemma.

Lemma 5.3 (Massart’s finite class lemma). *Let $(\sigma_1, \dots, \sigma_n)$ be n iid Rademacher random variables. For any finite subset $V \subseteq \mathbb{R}^n$ with $\|v\| \leq \sqrt{n}$ for all $v \in V$,*

$$\mathbb{E} \left[\max_{(v_1, \dots, v_n) \in V} \left| \frac{1}{n} \sum_{i=1}^n \sigma_i v_i \right| \right] \leq \sqrt{\frac{2 \ln(2|V|)}{n}}.$$

Proof. For each $v = (v_1, \dots, v_n) \in V' := V \cup (-V)$, let $\mathbf{x}_v := \frac{1}{n} \sum_{i=1}^n \sigma_i v_i$. Then

$$\max_{(v_1, \dots, v_n) \in V} \left| \frac{1}{n} \sum_{i=1}^n \sigma_i v_i \right| = \max_{v \in V'} \mathbf{x}_v.$$

For each $v = (v_1, \dots, v_n) \in V'$, the random variables $(v_1/n)\sigma_1, \dots, (v_n/n)\sigma_n$ are independent, and $(v_i/n)\sigma_i$ is $(v_i/n)^2$ -subgaussian (by Lemma 4.1), so $\mathbf{x}_v$ is $\|v\|^2/n^2$ -subgaussian:

$$K_{\mathbf{x}_v}(\lambda) \leq \frac{\|v\|^2 \lambda^2}{2n^2} \leq \frac{\lambda^2}{2n}.$$

Fix $\lambda > 0$. Then by Jensen's inequality,

$$\begin{aligned} \mathbb{E} \left[\max_{v \in V'} \mathbf{x}_v \right] &\leq \frac{1}{\lambda} \ln \mathbb{E} \left[\exp \left(\lambda \max_{v \in V'} \mathbf{x}_v \right) \right] \\ &= \frac{1}{\lambda} \ln \mathbb{E} \left[\max_{v \in V'} \exp(\lambda \mathbf{x}_v) \right] \\ &\leq \frac{1}{\lambda} \ln \mathbb{E} \left[\sum_{v \in V'} \exp(\lambda \mathbf{x}_v) \right] \\ &= \frac{1}{\lambda} \ln \left(\sum_{v \in V'} \exp(K_{\mathbf{x}_v}(\lambda)) \right). \end{aligned}$$

Now applying the bound on $K_{\mathbf{x}_v}(\lambda)$ for each $v \in V'$, we obtain

$$\mathbb{E} \left[\max_{v \in V'} \mathbf{x}_v \right] \leq \frac{\ln|V'|}{\lambda} + \frac{1}{\lambda} \cdot \frac{\lambda^2}{2n} = \frac{\ln|V'|}{\lambda} + \frac{\lambda}{2n}.$$

Plugging in $\lambda = \sqrt{2n \ln|V'|}$ and using the fact $|V'| \leq 2|V|$ gives the claim. $\square$

Lemma 5.3 gives the following bound on the conditional expectation from (5.6):

$$\mathbb{E} \left[\sup_{f \in \mathcal{F}} \left| \frac{1}{n} \sum_{i=1}^n \sigma_i f(\mathbf{z}_i) \right| \middle| \mathbf{z}_{1:n} \right] \leq \sqrt{\frac{2 \ln(2|\mathcal{F}[\{\mathbf{z}_1, \dots, \mathbf{z}_n\}])|}{n}}.$$

Now we appeal to Lemma 5.2 to bound $|\mathcal{F}[\{\mathbf{z}_1, \dots, \mathbf{z}_n\}]|$ by $\binom{n}{\leq d}$, where $d = \text{VCdim}(\mathcal{F})$. Putting this together with Proposition 5.7 proves Theorem 5.1.

Exercise 5.7. In the proof of Proposition 5.6, if we had taken expectations after (5.4), and then replaced the assignment $(\sigma_1, \dots, \sigma_n)$ with a random assignment $(\sigma_1, \dots, \sigma_n)$ (as in Proposition 5.7), we would have gotten the following variation on Proposition 5.7:

$$\mathbb{E} \left[\sup_{f \in \mathcal{F}} \left| \frac{1}{n} \sum_{i=1}^n f(\mathbf{z}_i) - \mathbb{E}[f(\mathbf{z})] \right| \right] = \mathbb{E} \left[\sup_{f \in \mathcal{F}} \left| \frac{1}{n} \sum_{i=1}^n \sigma_i (f(\mathbf{z}_i^a) - f(\mathbf{z}_i^b)) \right| \right].$$

- (a) Explain how to use Massart's finite class lemma (Lemma 5.3) and Sauer's lemma (Lemma 5.2) directly with the right-hand side of the above expression, without passing through Proposition 5.7. How does the result compare to Theorem 5.1?
- (b) Show that

$$\begin{aligned} \mathbb{E} \left[\sup_{f \in \mathcal{F}} \left| \frac{1}{n} \sum_{i=1}^n \sigma_i (f(\mathbf{z}_i^a) - f(\mathbf{z}_i^b)) \right| \right] &\geq \mathbb{E} \left[\sup_{f \in \mathcal{F}} \left| \frac{1}{n} \sum_{i=1}^n \sigma_i (f(\mathbf{z}_i) - \mathbb{E}[f(\mathbf{z})]) \right| \right] \\ &\geq \mathbb{E} \left[\sup_{f \in \mathcal{F}} \left| \frac{1}{n} \sum_{i=1}^n \sigma_i f(\mathbf{z}_i) \right| \right] - \frac{\sup_{f \in \mathcal{F}} |\mathbb{E}[f(\mathbf{z})]|}{\sqrt{n}}. \quad \diamond \end{aligned}$$

5.2.5 Zero-one loss class

How can Theorem 5.1 be used to understand the convergence of training error rates to population error rates for hypotheses in a hypothesis class $\mathcal{H} \subseteq \{0, 1\}^{\mathcal{X}}$? The key is simply to define a function class $\mathcal{F}_{\mathcal{H}} := \{f_h \mid h \in \mathcal{H}\}$ over the domain $\mathcal{Z} := \mathcal{X} \times \{0, 1\}$, where $f_h(x, y) = \mathbb{1}\{h(x) \neq y\}$. This function class $\mathcal{F}_{\mathcal{H}}$ is the *zero-one loss class* associated with $\mathcal{H}$. If P is a probability distribution on $\mathcal{Z} = \mathcal{X} \times \{0, 1\}$ and we consider an iid sample from P , then the population averages of $f_h \in \mathcal{F}_{\mathcal{H}}$ is the population error rate of h , and the empirical average of f_h is the training error rate of h .

Proposition 5.8. $\text{VCdim}(\mathcal{F}_{\mathcal{H}}) = \text{VCdim}(\mathcal{H})$.

Proof. We first show that $\text{VCdim}(\mathcal{F}_{\mathcal{H}}) \geq \text{VCdim}(\mathcal{H})$. Take any $T := \{x_1, \dots, x_d\} \subseteq \mathcal{X}$ shattered by $\mathcal{H}$, i.e., $\mathcal{H}[T] = \{0, 1\}^d$. Now define $T' := \{(x_1, 0), \dots, (x_d, 0)\} \subseteq \mathcal{X} \times \{0, 1\}$. For each $h \in \mathcal{H}$, we have $f_h(x_i, 0) = \mathbb{1}\{h(x_i) \neq 0\} = h(x_i)$. So $\mathcal{F}_{\mathcal{H}}[T'] = \mathcal{H}[T] = \{0, 1\}^d$, and hence T' is shattered by $\mathcal{F}_{\mathcal{H}}$. This implies $\text{VCdim}(\mathcal{F}_{\mathcal{H}}) \geq \text{VCdim}(\mathcal{H})$.

Now we show that $\text{VCdim}(\mathcal{F}_{\mathcal{H}}) \leq \text{VCdim}(\mathcal{H})$. Take any $T' := \{(x_1, y_1), \dots, (x_d, y_d)\} \subseteq \mathcal{X} \times \{0, 1\}$ of d distinct points shattered by $\mathcal{F}_{\mathcal{H}}$. We want to show that $T := \{x_1, \dots, x_d\}$ is shattered by $\mathcal{H}$.

We first show that $x_1, \dots, x_d$ are distinct points. To see this, suppose for sake of contradiction that we have $x_i = x_j$ for some $i \neq j$. Then it must be that $y_i \neq y_j$. But for any $h \in \mathcal{H}$, we have

$$f_h((x_i, y_i)) = \mathbb{1}\{h(x_i) \neq y_i\} = 1 - \mathbb{1}\{h(x_i) = y_i\} = 1 - \mathbb{1}\{h(x_j) = y_j\} = 1 - f_h((x_j, y_j)).$$

This means that no $f_h \in \mathcal{F}_{\mathcal{H}}$ can assign value 1 to both (x_i, y_i) and (x_j, y_j) , and hence T' is not shattered by $\mathcal{F}_{\mathcal{H}}$, a contradiction. So, $x_1, \dots, x_d$ are distinct, and $|T| = d$.

Now we show that T is shattered by $\mathcal{H}$. To see this, consider any labeling $(b_1, \dots, b_d) \in \{0, 1\}^d$. Since T is shattered by $\mathcal{F}_{\mathcal{H}}$, there exists $f_h \in \mathcal{F}_{\mathcal{H}}$ such that $f_h(x_i, y_i) = \mathbb{1}\{b_i \neq y_i\}$ for each $i \in [d]$. But for each $i \in [d]$, we have $f_h(x_i, y_i) = \mathbb{1}\{h(x_i) \neq y_i\}$. Therefore $h(x_i) = b_i$ for each $i \in [d]$. So T is shattered by $\mathcal{H}$. This implies $\text{VCdim}(\mathcal{F}_{\mathcal{H}}) \leq \text{VCdim}(\mathcal{H})$. $\square$

5.2.6 Lower bound for uniform convergence

Exercises 5.6 and 5.7 shows that Proposition 5.7 is tight up to constant factors and an additive $O(1/\sqrt{n})$ term. But what of the bound obtained via Massart's finite class lemma (Lemma 5.3)? This turns out to be almost tight—up to roughly a $\log(n)$ factor—for a worst-case choice of P .

Proposition 5.9. *For any $\mathcal{F} \subseteq \{0, 1\}^{\mathcal{Z}}$, there is a probability distribution P over $\mathcal{Z}$ such that, if $\mathbf{z}_{1:n}$ is a size- n iid sample from P , and $(\sigma_1, \dots, \sigma_n)$ are n iid Rademacher random variables regarded as independent of $\mathbf{z}_{1:n}$, then*

$$\mathbb{E} \left[\sup_{f \in \mathcal{F}} \left| \frac{1}{n} \sum_{i=1}^n \sigma_i f(\mathbf{z}_i) \right| \right] \geq \frac{1}{4} \min \left\{ 1, \sqrt{\frac{\text{VCdim}(\mathcal{F})}{n}} \right\}.$$

Proof. Let $d := \min\{\text{VCdim}(\mathcal{F}), n\}$. Let $T := \{t_1, \dots, t_d\}$ be a set of d points shattered by $\mathcal{F}$, and let P be the uniform distribution on T . Define $\mathbf{S}_j := \{i \in [n] \mid \mathbf{z}_i = t_j\}$ and $\mathbf{N}_j := |\mathbf{S}_j|$ for each $j \in [d]$. Then for each $f \in \mathcal{F}$,

$$\frac{1}{n} \sum_{i=1}^n \sigma_i f(\mathbf{z}_i) = \frac{1}{n} \sum_{j=1}^d \sum_{i \in \mathbf{S}_j} \sigma_i f(t_j).$$

Since T is shattered by $\mathcal{F}$, there exists $\mathbf{f}_o \in \mathcal{F}$ such that $\mathbf{f}_o(t_j) = \mathbb{1}\{\sum_{i \in \mathbf{S}_j} \sigma_i > 0\}$ for each $j \in [d]$. Therefore

$$\sup_{f \in \mathcal{F}} \left| \frac{1}{n} \sum_{i=1}^n \sigma_i f(\mathbf{z}_i) \right| \geq \left| \frac{1}{n} \sum_{j=1}^d \sum_{i \in \mathbf{S}_j} \sigma_i \mathbf{f}_o(t_j) \right| = \frac{1}{n} \sum_{j=1}^d \max \left\{ 0, \sum_{i \in \mathbf{S}_j} \sigma_i \right\},$$

and taking expectations,

$$\mathbb{E} \left[\sup_{f \in \mathcal{F}} \left| \frac{1}{n} \sum_{i=1}^n \sigma_i f(\mathbf{z}_i) \right| \right] \geq \frac{1}{n} \sum_{j=1}^d \mathbb{E} \left[\max \left\{ 0, \sum_{i \in \mathbf{S}_j} \sigma_i \right\} \right].$$

By the symmetry of the distribution of the Rademacher random variables, we have

$$\mathbb{E} \left[\max \left\{ 0, \sum_{i \in \mathbf{S}_j} \sigma_i \right\} \middle| \mathbf{z}_{1:n} \right] = \frac{1}{2} \mathbb{E} \left[\left| \sum_{i \in \mathbf{S}_j} \sigma_i \right| \middle| \mathbf{z}_{1:n} \right]$$

A tedious calculation (or Khintchine's inequality) shows that

$$\mathbb{E} \left[\left| \sum_{i \in \mathbf{S}_j} \sigma_i \right| \middle| \mathbf{z}_{1:n} \right] = \mathbf{N}_j \left(\frac{\mathbf{N}_j - 1}{\lfloor (\mathbf{N}_j - 1)/2 \rfloor} \right) 2^{1-\mathbf{N}_j} \geq \sqrt{\frac{\mathbf{N}_j}{2}}.$$

Therefore

$$\mathbb{E} \left[\sup_{f \in \mathcal{F}} \left| \frac{1}{n} \sum_{i=1}^n \sigma_i f(\mathbf{z}_i) \right| \right] \geq \frac{1}{2\sqrt{2}n} \mathbb{E} \left[\sum_{j=1}^d \sqrt{\mathbf{N}_j} \right].$$

By Hölder's inequality, we have

$$n = \sum_{j=1}^d \mathbf{N}_j = \sum_{j=1}^d \mathbf{N}_j^{1/3} \mathbf{N}_j^{2/3} \leq \left(\sum_{j=1}^d \mathbf{N}_j^{1/2} \right)^{2/3} \left(\sum_{j=1}^d \mathbf{N}_j^2 \right)^{1/3},$$

and hence

$$\sum_{j=1}^d \sqrt{\mathbf{N}_j} \geq \frac{n^{3/2}}{\sqrt{\sum_{j=1}^d \mathbf{N}_j^2}}.$$

By Jensen's inequality,

$$\mathbb{E} \left[\frac{n^{3/2}}{\sqrt{\sum_{j=1}^d \mathbf{N}_j^2}} \right] \geq \frac{n^{3/2}}{\sqrt{\mathbb{E} \left[\sum_{j=1}^d \mathbf{N}_j^2 \right]}} = \frac{n^{3/2}}{\sqrt{\sum_{j=1}^d \mathbb{E} \left[\mathbf{N}_j^2 \right]}}.$$

The random variables $(\mathbf{N}_1, \dots, \mathbf{N}_d)$ are the occupancies of d bins after throwing n balls independently into uniformly chosen bins, so

$$\mathbb{E} \left[\mathbf{N}_j^2 \right] = \text{var}(\mathbf{N}_j) + \mathbb{E}[\mathbf{N}_j]^2 = \frac{n}{d} \left(1 - \frac{1}{d} \right) + \left(\frac{n}{d} \right)^2 = \frac{n}{d} + \frac{n(n-1)}{d^2}.$$

We conclude

$$\mathbb{E} \left[\sup_{f \in \mathcal{F}} \left| \frac{1}{n} \sum_{i=1}^n \sigma_i f(\mathbf{z}_i) \right| \right] \geq \frac{1}{2\sqrt{2}n} \cdot \frac{n^{3/2}}{\sqrt{n + n(n-1)/d}} = \frac{1}{2\sqrt{2}} \sqrt{\frac{d}{d + n - 1}} \geq \frac{1}{4} \sqrt{\frac{d}{n}}. \quad \square$$

5.3 Empirical covering numbers (NOT DONE YET)

5.3.1 Detour: Pseudo-metric spaces

We say (X, d) is a metric space if X is a set of points and $d: X \times X \rightarrow \mathbb{R}$ is a metric, i.e., a function that satisfies the following:

1. $d(x, y) \geq 0$ for all $x, y \in X$, with equality if and only if $x = y$ (positivity);
2. $d(x, y) = d(y, x)$ for all $x, y \in X$ (symmetry);
3. $d(x, z) \leq d(x, y) + d(y, z)$ for all $x, y, z \in X$ (triangle inequality).

If d allows for distinct point $x \neq y$ to have $d(x, y) = 0$ but satisfies all other the properties required of a metric, then we say d is a pseudo-metric, and (X, d) is a pseudo-metric space.

Some examples of metric spaces are as follows.

- $(\mathbb{R}^n, \ell^p)$ where “ ℓ^p ” refers to the metric $d(x, y) = \|x - y\|_p$ based on the ℓ^p norm on $\mathbb{R}^n$.
- (V, d) where V is the vertex set of an undirected graph G that is connected and has only positive weight edges, and $d(u, v)$ is the length of the shortest path between u to v in G .
- $(\mathcal{H}, d)$ where $\mathcal{H} \subseteq \mathcal{Y}^{\mathcal{X}}$ is a hypothesis class, and $d(h, h') = \Pr(h(\mathbf{x}) \neq h'(\mathbf{x}))$ for a $\mathcal{X}$ -valued random variable $\mathbf{x}$ with $\Pr(\mathbf{x} = x) > 0$ for all $x \in \mathcal{X}$.

For the last example, if $\mathbf{x}$ is such that some distinct $h, h' \in \mathcal{H}$ have $\Pr(h(\mathbf{x}) = h'(\mathbf{x})) = 1$ (e.g., h and h' disagree only on some $x \in \mathcal{X}$ for which $\Pr(\mathbf{x} = x) = 0$), then d is not a metric, but it is a pseudo-metric.

For any $\epsilon > 0$, we say $C \subseteq X$ is an ϵ -cover of (X, d) if, for each $x \in X$, there exists $\tilde{x} \in C$ such that $d(x, \tilde{x}) \leq \epsilon$. Also, for any $\epsilon > 0$, we say $P \subseteq X$ is an ϵ -packing of (X, d) if $d(x, y) > \epsilon$ for all distinct $x, y \in P$.

Proposition 5.10. *Suppose (X, d) is a pseudo-metric space, and $P \subseteq X$ satisfies the following:*

- P is an ϵ -packing of (X, d)
- for all $x \in X \setminus P$, $P \cup \{x\}$ is not an ϵ -packing of (X, d) .

Then P is an ϵ -cover of (X, d) .

Proposition 5.10 shows that any maximal ϵ -packing is also an ϵ -cover. It suggests a simple “greedy” method for constructing an ϵ -cover:

- Start with $C := \emptyset$.
- While there exists $x \in X$ such that $d(x, y) > \epsilon$ for all $y \in C$:
 - Add such an $x \in X$ to C .

Proposition 5.11. *Suppose (X, d) is a pseudo-metric space, and for some $\epsilon > 0$, C is an ϵ -cover of (X, d) , and P is a 2ϵ -packing of (X, d) . Then $|P| \leq |C|$.*

The ϵ -covering number $\mathcal{N}(\epsilon, X, \mathbf{d})$ of $(X, \mathbf{d})$ is the smallest n such that there is an ϵ -cover of $(X, \mathbf{d})$ of cardinality n . The ϵ -packing number $\mathcal{M}(\epsilon, X, \mathbf{d})$ of $(X, \mathbf{d})$ is the largest m such that there is an ϵ -packing of $(X, \mathbf{d})$ of cardinality m . These numbers could be infinite, but Propositions 5.10 and 5.11 imply the following relationship between them.

Proposition 5.12. *For any pseudo-metric space $(X, \mathbf{d})$ and any $\epsilon > 0$,*

$$\mathcal{N}(\epsilon, X, \mathbf{d}) \leq \mathcal{M}(\epsilon, X, \mathbf{d}) \leq \mathcal{N}(\epsilon/2, X, \mathbf{d}).$$

Exercise 5.8. Prove Propositions 5.10 to 5.12. $\diamond$

Example 5.3. Let $X := [0, 1]^n$, and let $\mathbf{d}(x, y) := \|x - y\|_\infty$ is the ℓ^∞ distance. An ϵ -cover of $(X, \mathbf{d})$ is $C := \{0, \epsilon, 2\epsilon, \dots, 1 - \epsilon\}^n$, so $|C| = \lceil 1/\epsilon \rceil^n$. Hence $\mathcal{N}(\epsilon, [0, 1]^n, \ell^\infty) \leq \lceil 1/\epsilon \rceil^n$. $\diamond$

Example 5.4. Let $X := B^n$, where $B^n := \{x \in \mathbb{R}^n \mid \|x\| \leq 1\}$ is the unit ball in n -dimensional Euclidean space, and $\mathbf{d}(x, y) := \|x - y\|$ is the Euclidean distance (a.k.a. ℓ^2 -distance).

Suppose P is an ϵ -packing. Then balls of radius $\epsilon/2$ centered around points in P are disjoint. The total volume of these balls is

$$V_1 := |P|(\epsilon/2)^n v_n$$

where v_n is the volume of B^n itself. All of these balls are contained in a larger ball of radius $1 + \epsilon/2$, which has volume

$$V_2 := (1 + \epsilon/2)^n v_n.$$

Clearly, we must have $V_1 \leq V_2$, so

$$|P|(\epsilon/2)^n v_n \leq (1 + \epsilon/2)^n v_n,$$

and hence

$$|P| \leq (1 + 2/\epsilon)^n.$$

This implies $\mathcal{M}(\epsilon, B^n, \ell^2) \leq (1 + 2/\epsilon)^n$. By Proposition 5.12, we also have $\mathcal{N}(\epsilon, B^n, \ell^2) \leq (1 + 2/\epsilon)^n$.

Now suppose C is an ϵ -cover. Then B^n is contained in the union of $|C|$ balls of radius ϵ . By the union bound, the latter has volume at most $|C|\epsilon^n v_n$. So $|C| \geq (1/\epsilon)^n$. In particular, $\mathcal{N}(\epsilon, B^n, \ell^2) \geq (1/\epsilon)^n$. $\diamond$

5.3.2 Packing number bound in terms of VC dimension

The *disagreement (pseudo-)metric* $\mathbf{d}_P: \mathcal{F} \times \mathcal{F} \rightarrow \mathbb{R}$ on a function class $\mathcal{F} \subseteq \{0, 1\}^{\mathcal{Z}}$ corresponding to a probability distribution P over domain $\mathcal{Z}$ is given by

$$\mathbf{d}_P(f, f') := \Pr(f(\mathbf{z}) \neq f'(\mathbf{z})), \quad \mathbf{z} \sim P.$$

The disagreement metric is the same as that based on the $L^q(P)$ norm (up to a constant scaling) in the case $q = 1$; the norm for general $q \geq 1$ is given by

$$\|f\|_{L^q(P)} := \mathbb{E}[|f(\mathbf{z})|^q]^{1/q}, \quad \mathbf{z} \sim P.$$

For any $\epsilon > 0$, the ϵ -packing number—and also, via Proposition 5.12, the ϵ -covering number—for $(\mathcal{F}, \mathbf{d}_P)$ can be bounded in terms of ϵ and the VC dimension of $\mathcal{F}$, similar to that of the unit ball in $O(\text{VCdim}(\mathcal{F}))$ -dimensional Euclidean space.

Theorem 5.3. Let $\mathcal{F} \subseteq \{0, 1\}^{\mathcal{Z}}$ be a function class, and $\mathbf{d}_P$ be the disagreement metric on $\mathcal{F}$ corresponding to a probability distribution P on $\mathcal{Z}$. If $d := \text{VCdim}(\mathcal{F}) < \infty$, then for any $\epsilon \in (0, 1)$,

$$\mathcal{M}(\epsilon, \mathcal{F}, \mathbf{d}_P) \leq \left(\frac{4}{\epsilon}\right)^{2d}.$$

The proof of Theorem 5.3 is a simplified version of an argument due to Haussler (1995), who actually proves the much stronger bound of $e(d+1)(2e/\epsilon)^d$.

Proof of Theorem 5.3. Let $f_1, \dots, f_M \in \mathcal{F}$ be any functions from $\mathcal{F}$ satisfying $\mathbf{d}_P(f_i, f_j) > \epsilon$ for all $i \neq j$. Our goal is to upper-bound M . Let $(\mathbf{z}_1, \dots, \mathbf{z}_n) \sim P^n$ be a size- n iid sample from P , where

$$n := \frac{2 \ln(M)}{\epsilon}.$$

For any $i \neq j$,

$$\Pr(f_i(\mathbf{z}_t) = f_j(\mathbf{z}_t) \forall t \in [n]) = (1 - \mathbf{d}_P(f_i, f_j))^n < (1 - \epsilon)^n \leq \exp(-\epsilon n).$$

Therefore, by a union bound,

$$\Pr(\exists i \neq j \text{ s.t. } f_i(\mathbf{z}_t) = f_j(\mathbf{z}_t) \forall t \in [n]) \leq \binom{M}{2} \exp(-\epsilon n) < 1$$

where the last inequality follows by the choice of n . Therefore, there exists $z_{1:n} = (z_1, \dots, z_n) \in \mathcal{Z}^n$ such that the M functions $f_1, \dots, f_M$ exhibit distinct behaviors on $z_{1:n}$, which implies

$$|\mathcal{F}[\{z_1, \dots, z_n\}]| \geq M.$$

On the other hand, by Sauer's lemma (Lemma 5.2),

$$|\mathcal{F}[\{z_1, \dots, z_n\}]| \leq \binom{n}{\leq d}.$$

If $n \leq d$, then $\binom{n}{\leq d} = 2^n \leq 2^d$, so we conclude that $M \leq 2^d$. If $n > d$, then

$$M \leq \binom{n}{\leq d} \leq \left(\frac{ne}{d}\right)^d \leq \left(\frac{2e \ln(M)}{\epsilon d}\right)^d.$$

Observe that

$$\frac{2e \ln(M)}{\epsilon d} = \frac{4e}{\epsilon} \cdot \frac{\ln(M^{1/(2d)})}{M^{1/(2d)}} \cdot M^{1/(2d)} \leq \frac{4}{\epsilon} \cdot M^{1/(2d)}$$

since $\ln(x)/x \leq 1/e$ for all $x > 0$. Therefore

$$\left(\frac{2e \ln(M)}{\epsilon d}\right)^d \leq \left(\frac{4}{\epsilon}\right)^d \sqrt{M}.$$

We conclude that $M \leq (4/\epsilon)^{2d}$. □

Sauer's lemma (Lemma 5.2) can also be viewed as establishing a bound on the ϵ -covering number in the special case where P is the uniform distribution on a finite set of n points from $\mathcal{Z}$, and $\mathbf{d}$ is the $L^\infty(P)$ metric. It shows that $\mathcal{N}(\epsilon, \mathcal{F}, L^\infty(P)) \leq \binom{n}{\leq d}$ even as $\epsilon \rightarrow 0$.