## Solutions to Homework 5

## CS-526 Learning Theory

## Problem 1. VC dimension of union

1. Let $\mathcal{H}=\bigcup_{i=1}^{r} \mathcal{H}_{i}$. By definition of the growth function we have $\tau_{\mathcal{H}}(m) \leq \sum_{i=1}^{r} \tau_{\mathcal{H}_{i}}(m)$ for any set of $m$ points. If $k>d+1$ points are shattered by $\mathcal{H}$ then $2^{k}=\tau_{\mathcal{H}}(k) \leq$ $\sum_{i=1}^{r} \tau_{\mathcal{H}_{i}}(k) \leq r k^{d}$, where the last inequality follows directly from Sauer's lemma. Taking the logarithm on both sides and using the inequality yields

$$
k \leq \frac{4 d}{\log (2)} \log \left(\frac{2 d}{\log (2)}\right)+2 \frac{\log (r)}{\log (2)}
$$

Note that this inequality is trivially satisfied if $k \leq d+1$.
2. Assume that $k \geq 2 d+2$. It is enough to prove that $\tau_{\mathcal{H}_{1} \cup \mathcal{H}_{2}}(k)<2^{k}$.

$$
\begin{aligned}
& \tau_{\mathcal{H}_{1} \cup \mathcal{H}_{2}}(k) \leq \tau_{\mathcal{H}_{1}}(k)+\tau_{\mathcal{H}_{2}}(k) \leq \sum_{i=0}^{d}\binom{k}{i}+\sum_{i=0}^{d}\binom{k}{i}= \\
& =\sum_{i=0}^{d}\binom{k}{i}+\sum_{i=0}^{d}\binom{k}{k-i}=\sum_{i=0}^{d}\binom{k}{i}+\sum_{i=k-d}^{k}\binom{k}{i} \leq \\
& \leq \sum_{i=0}^{d}\binom{k}{i}+\sum_{i=d+2}^{k}\binom{k}{i}<\sum_{i=0}^{d}\binom{k}{i}+\sum_{i=d+1}^{k}\binom{k}{i}= \\
& =\sum_{i=0}^{k}\binom{k}{i}=2^{k}
\end{aligned}
$$

Lemma (Sauer-Shelah-Perles) Let $\mathcal{H}$ be a hypothesis class with $\operatorname{VCdim}(H) \leq d<$ $\infty$ and growth function $\tau_{\mathcal{H}}$. Then, for all $m, \tau_{\mathcal{H}}(m) \leq \sum_{i=0}^{d}\binom{m}{i}$. In particular, if $m>d+1$ and $d>2$ then $\tau_{\mathcal{H}}(m)<m^{d}$.

## Problem 2. Least squares and regularized least squares

1. We have

$$
\begin{gathered}
\hat{\beta}=\underset{\beta \in \mathbb{R}^{d}}{\arg \min } \mathcal{J}(\beta):=\|y-X \beta\|^{2} \\
\mathcal{J}(\beta)=\beta^{T} X^{T} X \beta-2 \beta^{T} X^{T} y+y^{T} y \\
\nabla \mathcal{J}(\beta)=2\left(X^{T} X \beta-X^{T} y\right)
\end{gathered}
$$

Equating $\nabla \mathcal{J}$ to 0 , we get $\hat{\beta}=\left(X^{T} X\right)^{-1} X^{T} y$.
2. In this case, we have

$$
\begin{gathered}
\hat{\beta}=\underset{\beta \in \mathbb{R}^{d}}{\arg \min } \mathcal{J}^{\prime}(\beta):=\|y-X \beta\|^{2}+\lambda\|\beta\|^{2} \\
\mathcal{J}^{\prime}(\beta)=\beta^{T} X^{T} X \beta-2 \beta^{T} X^{T} y+y^{T} y+\lambda \beta^{T} \lambda \\
\nabla \mathcal{J}^{\prime}(\beta)=2\left(X^{T} X \beta-X^{T} y+\lambda \beta\right)
\end{gathered}
$$

Equating $\nabla \mathcal{J}$ to 0 , we get $\hat{\beta}=\left(X^{T} X+\lambda I_{d}\right)^{-1} X^{T} y$.
Increasing the regularization parameter reduces the variance of the model at the cost of increasing its bias towards solutions with a small $l_{2}$-norm.

## Problem 3. Linear regression with projections

Refer to the lecture notes.

## Problem 4. Bias-variance decomposition

The three contributions are

$$
\begin{aligned}
\text { Noise } & =\mathbb{E}_{x, y}\left[(\bar{h}(x)-y)^{2}\right] \\
(\text { Bias })^{2} & =\mathbb{E}_{x}\left[\left(\mathbb{E}_{S}\left[h_{S}(x)\right]-\bar{h}(x)\right)^{2}\right] \\
\text { Variance } & =\mathbb{E}_{S} \mathbb{E}_{x \mid S}\left[\left(h_{S}(x)-\mathbb{E}_{S}\left[h_{S}(x)\right]\right)^{2}\right] .
\end{aligned}
$$

First let us compute the optimal estimator $\bar{h}$ :

$$
\bar{h}(x)=\mathbb{E}[y \mid x]=\beta^{T} x
$$

With this we can already compute the noise part:

$$
\mathbb{E}_{x, y}\left[(\bar{h}(x)-y)^{2}\right]=\mu^{2} \mathbb{E}\left[\epsilon^{2}\right]=\mu^{2}
$$

Let's now focus on the data-dependent estimator:

$$
h_{S}(x)= \begin{cases}\left(\left(X_{\mathcal{A}}^{T} X_{\mathcal{A}}\right)^{-1} X_{\mathcal{A}}^{T} y\right)^{T} x_{\mathcal{A}}, & p<n-1 \\ \left(X_{\mathcal{A}}^{T}\left(X_{\mathcal{A}}^{T} X_{\mathcal{A}}\right)^{\dagger} y\right)^{T} x_{\mathcal{A}}, & p>n+1\end{cases}
$$

Consider the quantity $\mathbb{E}\left[h_{S}(x)\right]$ for $p<n-1$ :

$$
\begin{aligned}
\mathbb{E}\left[\left(\left(X_{\mathcal{A}}^{T} X_{\mathcal{A}}\right)^{-1} X_{\mathcal{A}}^{T} y\right)^{T} x_{\mathcal{A}}\right] & =\mathbb{E}\left[\left(X_{\mathcal{A}} \beta_{\mathcal{A}}+X_{\mathcal{A}^{C}} \beta_{\mathcal{A}^{C}}+\mu \epsilon\right)^{T} X_{\mathcal{A}}\left(X_{\mathcal{A}}^{T} X_{\mathcal{A}}\right)^{-1} x_{\mathcal{A}}\right] \\
& =\mathbb{E}\left[\beta_{\mathcal{A}}^{T} X_{\mathcal{A}}^{T} X_{\mathcal{A}}\left(X_{\mathcal{A}}^{T} X_{\mathcal{A}}\right)^{-1} x_{\mathcal{A}}\right] \\
& =\beta_{\mathcal{A}}^{T} x_{\mathcal{A}} .
\end{aligned}
$$

Similarly for $p>n+1$ :

$$
\begin{aligned}
\mathbb{E}\left[\left(X_{\mathcal{A}}^{T}\left(X_{\mathcal{A}}^{T} X_{\mathcal{A}}\right)^{\dagger} y\right)^{T} x_{\mathcal{A}}\right] & =\beta_{\mathcal{A}}^{T} \mathbb{E}\left[X_{\mathcal{A}}^{T}\left(X_{\mathcal{A}} X_{\mathcal{A}}^{T}\right)^{\dagger} X_{\mathcal{A}}\right] x_{\mathcal{A}} \\
& =\frac{n}{p} \beta_{\mathcal{A}} x_{\mathcal{A}} .
\end{aligned}
$$

Let's define the following quantity:

$$
\psi= \begin{cases}1 ; & p<n-1 \\ n / p ; & p>n+1\end{cases}
$$

Then we have

$$
\mathbb{E}\left[h_{S}(x)\right]=\psi \beta_{\mathcal{A}} x_{\mathcal{A}} .
$$

Computation of the (Bias) ${ }^{2}$ contribution:

$$
\begin{aligned}
\mathbb{E}_{x}\left[\left(\mathbb{E}_{S}\left[h_{S}(x)\right]-\bar{h}(x)\right)^{2}\right] & =\mathbb{E}_{x}\left[\left(\psi \beta_{\mathcal{A}}^{T} x_{\mathcal{A}}-\beta^{T} x\right)^{2}\right] \\
& =\mathbb{E}_{x}\left[\left((\psi-1) \beta_{\mathcal{A}}^{T} x_{\mathcal{A}}-\beta_{\mathcal{A}^{C}}^{T} x_{\mathcal{A}^{C}}\right)^{2}\right] \\
& =\mathbb{E}_{x}\left[(\psi-1)^{2}\left(\beta_{\mathcal{A}}^{T} x_{\mathcal{A}}\right)^{2}\right]+\mathbb{E}_{x}\left[\left(\beta_{\mathcal{A}^{C}}^{T} x_{\mathcal{A}^{C}}\right)^{2}\right] \\
& =(\psi-1)^{2}\left\|\beta_{\mathcal{A}}\right\|^{2}+\left\|\beta_{\mathcal{A}^{C}}\right\|^{2} .
\end{aligned}
$$

We now compute the variance:

$$
\begin{aligned}
\mathbb{E}_{S} \mathbb{E}_{x \mid S}\left[\left(h_{S}(x)-\mathbb{E}_{S}\left[h_{S}(x)\right]\right)^{2}\right] & =\mathbb{E}_{S} \mathbb{E}_{x \mid S}\left[\left(\hat{\beta}_{\mathcal{A}} x_{\mathcal{A}}-\psi \beta_{\mathcal{A}} x_{\mathcal{A}}\right)^{2}\right] \\
& =\mathbb{E}_{S}\left[\left\|\hat{\beta}_{\mathcal{A}}-\psi \beta_{\mathcal{A}}\right\|^{2}\right]=\mathbb{E}_{S}\left[\left\|\beta_{\mathcal{A}}-\hat{\beta}_{\mathcal{A}}+(\psi-1) \beta_{\mathcal{A}}\right\|^{2}\right] \\
& =\mathbb{E}_{S}\left[\left\|\beta_{\mathcal{A}}-\hat{\beta}_{\mathcal{A}}\right\|^{2}+\left(\psi^{2}-2 \psi+1\right)\left\|\beta_{\mathcal{A}}\right\|^{2}+2\left(\beta_{\mathcal{A}}-\hat{\beta}_{\mathcal{A}}\right)^{T} \beta_{\mathcal{A}}(\psi-1)\right] \\
& =\mathbb{E}_{S}\left[\left\|\beta_{\mathcal{A}}-\hat{\beta}_{\mathcal{A}}\right\|^{2}\right]+\left(\psi^{2}-1\right)\left\|\beta_{\mathcal{A}}\right\|^{2}-2(\psi-1) \mathbb{E}_{S}\left[\beta_{\mathcal{A}}^{T} \hat{\beta}_{\mathcal{A}}\right] .
\end{aligned}
$$

Focusing on the last term which exists only when $p>n+1$, we have:

$$
\begin{aligned}
\mathbb{E}_{S}\left[\beta_{\mathcal{A}}^{T} \hat{\beta}_{\mathcal{A}}\right] & =\mathbb{E}_{S}\left[\beta_{\mathcal{A}}^{T} X_{\mathcal{A}}^{T}\left(X_{\mathcal{A}} X_{\mathcal{A}}^{T}\right)^{\dagger} y\right] \\
& =\mathbb{E}_{S}\left[\beta_{\mathcal{A}}^{T} X_{\mathcal{A}}^{T}\left(X_{\mathcal{A}} X_{\mathcal{A}}^{T}\right)^{\dagger}\left(X_{\mathcal{A}} \beta_{\mathcal{A}}+X_{\mathcal{A}^{c}} \beta_{\mathcal{A}^{C}}+\mu \epsilon\right)\right] \\
& =\mathbb{E}_{S}\left[\beta_{\mathcal{A}}^{T} X_{\mathcal{A}}^{T}\left(X_{\mathcal{A}} X_{\mathcal{A}}^{T}\right)^{\dagger} X_{\mathcal{A}} \beta_{\mathcal{A}}\right] \\
& =\mathbb{E}_{S}\left[\operatorname{Tr}\left\{\beta_{\mathcal{A}} \beta_{\mathcal{A}}^{T} X_{\mathcal{A}}^{T}\left(X_{\mathcal{A}} X_{\mathcal{A}}^{T}\right)^{\dagger} X_{\mathcal{A}}\right\}\right] \\
& =\operatorname{Tr}\left\{\beta_{\mathcal{A}} \beta_{\mathcal{A}}^{T} \mathbb{E}_{S}\left[X_{\mathcal{A}}^{T}\left(X_{\mathcal{A}} X_{\mathcal{A}}^{T}\right)^{\dagger} X_{\mathcal{A}}\right]\right\} \\
& =\operatorname{Tr}\left\{\beta_{\mathcal{A}} \beta_{\mathcal{A}}^{T} I_{p} \frac{n}{p}\right\}=\frac{n}{p}\left\|\beta_{\mathcal{A}}\right\|^{2} .
\end{aligned}
$$

Plugging back, we get

$$
\mathbb{E}_{S} \mathbb{E}_{x \mid S}\left[\left(h_{S}(x)-\mathbb{E}_{S}\left[h_{S}(x)\right]\right)^{2}\right]=\mathbb{E}_{S}\left[\left\|\beta_{\mathcal{A}}-\hat{\beta}_{\mathcal{A}}\right\|^{2}\right]-(1-\psi)^{2}\left\|\beta_{\mathcal{A}}\right\|^{2} .
$$

Therefore,

$$
\text { Variance }= \begin{cases}\mathbb{E}_{S}\left[\left\|\beta_{\mathcal{A}}-\hat{\beta}_{\mathcal{A}}\right\|^{2}\right], & p<n-1 \\ \mathbb{E}_{S}\left[\left\|\beta_{\mathcal{A}}-\hat{\beta}_{\mathcal{A}}\right\|^{2}\right]-\left(1-\frac{n}{p}\right)^{2}\left\|\beta_{\mathcal{A}}\right\|^{2}, & p>n+1\end{cases}
$$

By using the expression obtained for $\mathbb{E}_{S}\left[\left\|\beta_{\mathcal{A}}-\hat{\beta}_{\mathcal{A}}\right\|^{2}\right]$ in the class, we have the following contributions to the error.
For $p<n-1$ :

$$
\text { Error }=\underbrace{\mu^{2}}_{\text {Noise }}+\underbrace{\left\|\beta_{\mathcal{A}^{c}}\right\|^{2}}_{\text {Bias }^{2}}+\underbrace{\frac{p}{n-p-1}\left(\mu^{2}+\left\|\beta_{\mathcal{A}^{c}}\right\|^{2}\right)}_{\text {Variance }}
$$

For $p>n+1$ :

$$
\begin{aligned}
& \text { Error }=\underbrace{\mu^{2}}_{\text {Noise }}+\underbrace{\left\|\beta_{\mathcal{A}^{C}}\right\|^{2}+(1-n / p)^{2}\left\|\beta_{\mathcal{A}}\right\|^{2}}_{\text {Bias }^{2}} \\
&+\underbrace{(1-n / p)\left\|\beta_{\mathcal{A}^{C}}\right\|^{2}+\frac{n}{p-n-1}\left(\mu^{2}+\left\|\beta_{\mathcal{A}^{c}}\right\|^{2}\right)-(1-n / p)^{2}\left\|\beta_{\mathcal{A}}\right\|^{2}}_{\text {Variance }}
\end{aligned}
$$

Define $\alpha=p / n$ and $\varphi=n / d$. Assume $\mathcal{A}$ as a uniformly random subset of $1,2, \cdots, d$ and taking $p, n, d \rightarrow \infty$ with $\alpha$ and $\varphi$ finite. For $\alpha<1$ :

$$
\text { Error }=\underbrace{\mu^{2}}_{\text {Noise }}+\underbrace{(1-\alpha \varphi)\|\beta\|^{2}}_{\text {Bias }^{2}}+\underbrace{\mu^{2}+(1-\alpha \varphi)\|\beta\|^{2} \frac{\alpha}{1-\alpha}}_{\text {Variance }}
$$

For $\alpha>1$ :

$$
\begin{aligned}
& \text { Error }=\underbrace{\mu^{2}}_{\text {Noise }}+\underbrace{(1-\alpha \varphi)\|\beta\|^{2}+}_{\text {Bias }^{2}}+\frac{(\alpha-1)^{2}}{\alpha} \varphi\|\beta\|^{2} \\
&+\underbrace{(\alpha-1) \varphi\|\beta\|^{2}+\frac{\mu^{2}+(1-\alpha \varphi)}{\alpha-1}\|\beta\|^{2}-\frac{(\alpha-1)^{2}}{\alpha} \varphi\|\beta\|^{2}}_{\text {Variance }}
\end{aligned}
$$



