

Homework 3: Bias and Variance Trade-off

Question 1. Ensemble and Bias-Variance Trade-off As a reminder, the expected generalization error enjoys the following *bias-variance* trade-off:

$$\mathbb{E} \left[|h(x; \mathcal{D}) - y|^2 \right] = \underbrace{\mathbb{E} \left[(\mathbb{E}[h(x; \mathcal{D})|x] - y_*(x))^2 \right]}_{\textcircled{2} \text{ bias}} + \underbrace{\mathbb{E} \left[(h(x; \mathcal{D}) - \mathbb{E}[h(x; \mathcal{D})|x])^2 \right]}_{\textcircled{3} \text{ variance}} + \underbrace{\mathbb{E} \left[(y_*(x) - y)^2 \right]}_{\text{Irreducible error}},$$

where x, y represent the sampled test data. \mathcal{D} represents the sampled training dataset. $h(\cdot; \mathcal{D})$ is our prediction model learned using this dataset (i.e. $h(\cdot; \mathcal{D})$ is the learnt hypothesis given the training examples). $y_*(x)$ is the true model we want to learn and $\mathbb{E}[y|x] = y_*(x)$. In the following questions, we are interested in the generalization performance of ensemble.

1.0.1. Weight Average or Prediction Average? Does the ensemble of linear models using weight average or prediction average give the same expected generalization error? Does the ensemble of (nonlinear) neural networks using weight average or prediction average give the same expected generalization error?

1.0.2. Bagging - Uncorrelated Models. One way to construct an ensemble is through bootstrap aggregation, or bagging, that takes a dataset \mathcal{D} and generates k new datasets, with replacement. In this question, we assume the generated dataset \mathcal{D}_i has the same size as \mathcal{D} . Then train a model for each dataset, \mathcal{D}_i , resulting in k models. The ensemble model is following:

$$\bar{h}(x; \mathcal{D}) = \frac{1}{k} \sum_{i=1}^k h(x; \mathcal{D}_i)$$

For this part, we will make a very unrealistic assumption that the predictions of the ensemble members are uncorrelated. That is $\text{Cov}(h(x; \mathcal{D}_j), h(x; \mathcal{D}_k)) = 0$.

1.0.2.1. Bias with bagging. Show that ensemble does not change the bias term in the generalization error.

$$\text{Show bias} = \mathbb{E} \left[(\mathbb{E}[\bar{h}(x; \mathcal{D})|x] - y_*(x))^2 \right] = \mathbb{E} \left[(\mathbb{E}[h(x; \mathcal{D})|x] - y_*(x))^2 \right]$$

1.0.2.2. Variance with bagging. Assume the variance of a single predictor is σ^2 , $\mathbb{E}[(h(x; \mathcal{D}) - \mathbb{E}[h(x; \mathcal{D})|x])^2] = \sigma^2$. Derive the variance of ensemble in terms of σ^2 under uncorrelated predictions.

1.0.3. Bagging - General Case. In practice, there will be correlations among the k predictions of the ensemble members because the sampled training datasets would be very similar to each other. For simplicity, assume a non-zero pairwise correlation between the ensemble members, that is, ρ . The variance of the predictor for $h(x; \mathcal{D}_j)$ is σ_j^2 .

$$\rho = \frac{\text{Cov}(h(x; \mathcal{D}_j), h(x; \mathcal{D}_k))}{\sigma_j \sigma_k} \quad \forall j \neq k$$

1.0.3.1. Bias under Correlation. Does the correlation change the bias term in the generalization error? If so, derive the new expression in terms of ρ . Provide your justification.

1.0.3.2. Variance under Correlation. Assume the variance of a single predictor is σ^2 . Derive the variance term of ensemble in terms of σ and ρ .

$$\text{Show variance} = \mathbb{E} \left[(\bar{h}(x; \mathcal{D}) - \mathbb{E}[\bar{h}(x; \mathcal{D})|x])^2 \right] = \left(\rho + \frac{1 - \rho}{k} \right) \sigma^2$$

1.0.3.3. Intuitions on bagging. Based on the derived variance, what happens to the variance when you increase the number of ensemble models k ? What do $\rho = 0$, $\rho = 1$ represent and their consequences for the variance?

Question 2. (Central Limit Theorem for Kernel Density Estimator) Let \hat{p}_h be the kernel density estimator (in one dimension) with bandwidth $h = h_n$. Let $s_n^2(x) = \text{Var}(\hat{p}_h(x))$.

(a) Show that

$$\frac{\hat{p}_h(x) - p(x)}{s_n(x)} \xrightarrow{d} N(0, 1)$$

where $p_h(x) = \mathbb{E}[\hat{p}_h(x)]$.

Hint: Recall that the Lyapunov central limit theorem says the following: Suppose that Y_1, Y_2, \dots are independent. Let $\mu_i = \mathbb{E}[Y_i]$ and $\sigma_i^2 = \text{Var}(Y_i)$. Let $s_n^2 = \sum_{i=1}^n \sigma_i^2$. If

$$\lim_{n \rightarrow \infty} \frac{1}{s_n^{2+\delta}} \sum_{i=1}^n \mathbb{E}[|Y_i - \mu_i|^{2+\delta}] = 0$$

for some $\delta > 0$, then $s_n^{-1} \sum_{i=1}^n (Y_i - \mu_i) \xrightarrow{d} N(0, 1)$.

The we need to show the following two statements

- Let $p > 1$, we can show $\mathbb{E} \left[\left| \frac{1}{h} K \left(\frac{\|x - X_i\|}{h} \right) - p_h(x) \right|^p \right] = \Theta \left(\frac{1}{h^{p-1}} \right)$ using $\frac{1}{2^p} |a|^p - |b|^p \leq |a - b|^p \leq 2^p |a|^p + 2^p |b|^p$, $\mathbb{E}[|Z_i|^p] = \frac{1}{h^{p-1}} \int |K(\|v\|)|^p p(x + hv) dv$ and $|K(\|v\|)|^p p(x + hv) dv \rightarrow |K(\|v\|)|^p p(x + hv) dv$.
- $s_n^2 = \frac{1}{n} \mathbb{E} \left[\left| \frac{1}{h} K \left(\frac{\|x - X_i\|}{h} \right) - p_h(x) \right|^2 \right] = \Theta \left(\frac{1}{nh} \right)$.

(b) Assume that the smoothness is $\beta = 2$. Suppose that the bandwidth h_n is chosen optimally. Show that

$$\frac{\hat{p}_h(x) - p(x)}{s_n(x)} \xrightarrow{d} N(b(x), 1)$$

for some constant $b(x)$ which is, in general, not 0.

Question 3. Estimating the Sobolev Ellipsoid via Spectral Methods Let $X_1, \dots, X_n \sim P$ where $X_i \in [0, 1]$ and P has density p . Let $\varphi_1, \varphi_2, \dots$ be an orthonormal basis for $L_2[0, 1]$. Hence $\int_0^1 \varphi_j^2(x) dx = 1$ for all j and $\int_0^1 \varphi_j(x) \varphi_k(x) dx = 0$ for $j \neq k$. Assume that the basis is uniformly bounded, i.e. $\sup_j \sup_{0 \leq x \leq 1} |\varphi_j(x)| \leq C < \infty$. We may expand p as $p(x) = \sum_{j=1}^{\infty} \beta_j \varphi_j(x)$ where $\beta_j = \int \varphi_j(x) p(x) dx$. Define

$$\hat{p}(x) = \sum_{j=1}^k \hat{\beta}_j \varphi_j(x)$$

where $\hat{\beta}_j = (1/n) \sum_{i=1}^n \varphi_j(X_i)$.

(a) Show that the risk is bounded by

$$\frac{ck}{n} + \sum_{j=k+1}^{\infty} \beta_j^2$$

for some constant $c > 0$.

(b) Define the Sobolev ellipsoid $E(m, L)$ of order m as the set of densities of the form $p(x) = \sum_{j=1}^{\infty} \beta_j \varphi_j(x)$ where $\sum_{j=1}^{\infty} \beta_j^2 j^{2m} < L^2$. Show that the risk for any density in $E(m, L)$ is bounded by $c[(k/n) + (1/k)^{2m}]$. Using this bound, find the optimal value of k and find the corresponding risk.