## STAT 535: Statistical Machine Learning

## Lecture 3: Regression: Nonparametric approaches

Instructor: Yen-Chi Chen

Reference: Section 6 of All of Nonparametric Statistics by Larry Wasserman.

### 3.1 Introduction

Let $\left(X_{1}, Y_{1}\right), \cdots,\left(X_{n}, Y_{n}\right)$ be a bivariate random sample. In the regression analysis, we are often interested in the regression function

$$
m(x)=\mathbb{E}(Y \mid X=x) .
$$

Sometimes, we will write

$$
Y_{i}=m\left(X_{i}\right)+\epsilon_{i},
$$

where $\epsilon_{i}$ is a mean 0 noise. The simple linear regression model is to assume that $m(x)=\beta_{0}+\beta_{1} x$, where $\beta_{0}$ and $\beta_{1}$ are the intercept and slope parameter. In the first part of the lecture, we will talk about methods that direct estimate the regression function $m(x)$ without imposing any parametric form of $m(x)$. This approach is called the nonparametric regression.

### 3.2 Regressogram (Binning)

We start with a very simple but extremely popular method. This method is called regressogram but people often call it binning approach. You can view it as

$$
\text { regressogram }=\text { regression }+ \text { histogram } .
$$

For simplicity, we assume that the covariates $X_{i}$ 's are from a distribution over $[0,1]$.
Similar to the histogram, we first choose $M$, the number of bins. Then we partition the interval $[0,1]$ into $M$ equal-width bins:

$$
B_{1}=\left[0, \frac{1}{M}\right), B_{2}=\left[\frac{1}{M}, \frac{2}{M}\right), \cdots, B_{M-1}=\left[\frac{M-2}{M}, \frac{M-1}{M}\right), B_{M}=\left[\frac{M-1}{M}, 1\right] .
$$

When $x \in B_{\ell}$, we estimate $m(x)$ by

$$
\widehat{m}_{M}(x)=\frac{\sum_{i=1}^{n} Y_{i} I\left(X_{i} \in B_{\ell}\right)}{\sum_{i=1}^{n} I\left(X_{i} \in B_{\ell}\right)}=\text { average of the responses whose covariates is in the same bin as } x .
$$

Theorem 3.1 Assume that the PDF of $X p(x) \geq p_{0}>0$ for all $x \in[0,1]$ and $\mathbb{E}\left(Y^{2} \mid X=x\right)<\infty$. Then

$$
\operatorname{bias}\left(\widehat{m}_{M}(x)\right)=O\left(\frac{1}{M}\right), \quad \operatorname{Var}\left(\widehat{m}_{M}(x)\right)=O\left(\frac{M}{n}\right) .
$$

Proof: Suppose that $x$ belongs to bin $B_{\ell}$. Let $\mu_{M}(x)=\mathbb{E}\left(Y_{i} I\left(X_{i} \in B_{\ell}\right)\right)$ and $q_{M}(x)=\mathbb{E}\left(I\left(X_{i} \in B_{\ell}\right)\right)$ and let $\mu(x)=m(x) \cdot p(x)$.

Using $\widehat{m}_{M}(x)=\frac{\widehat{\mu}_{M}(x)}{\widehat{q}_{M}(x)}$, where

$$
\widehat{\mu}_{M}(x)=\frac{1}{n} \sum_{i=1}^{n} Y_{i} I\left(X_{i} \in B_{\ell}\right), \quad \widehat{q}_{M}(x)=\sum_{i=1}^{n} I\left(X_{i} \in B_{\ell}\right)
$$

the difference can be decomposed into

$$
\widehat{m}_{M}(x)-m(x)=\underbrace{\frac{\widehat{\mu}_{M}(x)}{\widehat{q}_{M}(x)}-\frac{\mu_{M}(x)}{q_{M}(x)}}_{\sim \text { Variance }}+\underbrace{\frac{\mu_{M}(x)}{q_{M}(x)}-\frac{\mu(x)}{p(x)}}_{\text {bias }} .
$$

Bias. Since we have $M$ bins, the width of each bin is $1 / M$. Thus, it is easy to see that $M \cdot q_{M}(x)$ can be viewed as a density histogram estimator of $p(x)$. Therefore, by the theory of histogram, the bias will be $p(x)-M \cdot q_{M}(x)=O(1 / M)$. Similarly, one can show that $M \cdot \mu_{M}(x)$ can be viewed as an estimator of $\mu(x)$ and $\mu(x)-M \cdot \mu_{M}(x)=O(1 / M)$. Using the fact that $\frac{1}{1+\epsilon}=1-\epsilon+O\left(\epsilon^{2}\right)$ when $\epsilon \rightarrow 0$, we conclude that the bias part

$$
\begin{aligned}
\frac{\mu_{M}(x)}{q_{M}(x)}-\frac{\mu(x)}{p(x)} & =\frac{M \mu_{M}(x)}{M q_{M}(x)}-\frac{\mu(x)}{p(x)} \\
& =\frac{\mu(x)+O(1 / M)}{p(x)+O(1 / M)}-\frac{\mu(x)}{p(x)} \\
& =\frac{O(1 / M)}{p(x)}+\frac{\mu(x)}{p^{2}(x)} O(1 / M) \\
& =O(1 / M)
\end{aligned}
$$

Variance. For the variance part, it is easy to see that $\mathbb{E}\left(\widehat{\mu}_{M}(x)\right)=\mu_{M}(x)$ and $\mathbb{E}\left(\widehat{q}_{M}(x)\right)=q_{M}(x)$. Also, it is easy to see that the variance (using the same derivation as histogram),

$$
\operatorname{Var}\left(\widehat{\mu}_{M}(x)\right)=O\left(\frac{1}{M n}\right), \quad \operatorname{Var}\left(\widehat{q}_{M}(x)\right)=O\left(\frac{1}{M n}\right)
$$

Thus,

$$
\operatorname{Var}\left(M \widehat{\mu}_{M}(x)\right)=O\left(\frac{M}{n}\right), \quad \operatorname{Var}\left(M \widehat{q}_{M}(x)\right)=O\left(\frac{M}{n}\right)
$$

Let $\Delta_{\mu}(x)=M \widehat{\mu}_{M}(x)-M \mu_{M}(x)=O_{P}(\sqrt{M / n})$ and $\Delta_{q}(x)=M \widehat{q}_{M}(x)-M q_{M}(x)=O_{P}(\sqrt{M / n})$. Then

$$
\begin{aligned}
\frac{\widehat{\mu}_{M}(x)}{\widehat{q}_{M}(x)}-\frac{\mu_{M}(x)}{q_{M}(x)} & =\frac{M \widehat{\mu}_{M}(x)}{M \widehat{q}_{M}(x)}-\frac{M \mu_{M}(x)}{M q_{M}(x)} \\
& =\frac{M \mu_{M}(x)+\Delta_{\mu}(x)}{M q_{M}(x)+\Delta_{q}(x)}-\frac{M \mu_{M}(x)}{M q_{M}(x)} \\
& =\frac{M \mu_{M}(x)+\Delta_{\mu}(x)}{M q_{M}(x)}-\frac{M \mu_{M}(x)}{M^{2} q_{M}^{2}(x)} \Delta_{q}(x)-\frac{M \mu_{M}(x)}{M q_{M}(x)}+\text { smaller order terms } \\
& \approx \frac{1}{M q_{M}(x)} \Delta_{\mu}(x)-\frac{M \mu_{M}(x)}{M^{2} q_{M}^{2}(x)} \Delta_{q}(x) \\
& \approx \frac{1}{p(x)} \Delta_{\mu}(x)-\frac{\mu(x)}{p^{2}(x)} \Delta_{q}(x)
\end{aligned}
$$

Note that the last $\approx \operatorname{sign}$ is due to the bias analysis. Thus, the variance of $\frac{\widehat{\mu}_{M}(x)}{\widehat{q}_{M}(x)}$ equals the variance of $\frac{1}{p(x)} \Delta_{\mu}(x)-\frac{\mu(x)}{p^{2}(x)} \Delta_{q}(x)$, which is of rate $O(M / n)$.

Therefore, the MSE and MISE will be at rate

$$
\mathbf{M S E}=O\left(\frac{1}{M^{2}}\right)+O\left(\frac{M}{n}\right), \quad \text { MISE }=O\left(\frac{1}{M^{2}}\right)+O\left(\frac{M}{n}\right)
$$

leading to the optimal number of bins $M^{*} \asymp n^{1 / 3}$ and the optimal convergence rate $O\left(n^{-2 / 3}\right)$, the same as the histogram.

Similar to the histogram, the regressogram has a slower convergence rate compared to many other competitors (we will introduce several other candidates). However, they (histogram and regressogram) are still very popular because the construction of an estimator is very simple and intuitive; practitioners with little mathematical training can easily master these approaches.

Note that if we assume that the response variable $Y$ is bounded, you can construct a similar concentration bound as the case of histogram and obtain the rate under the $L_{\infty}$ metric.

### 3.3 Kernel Regression

Given a point $x_{0}$, assume that we are interested in the value $m\left(x_{0}\right)$. Here is a simple method to estimate that value. When $m\left(x_{0}\right)$ is smooth, an observation $X_{i} \approx x_{0}$ implies $m\left(X_{i}\right) \approx m\left(x_{0}\right)$. Thus, the response value $Y_{i}=m\left(X_{i}\right)+\epsilon_{i} \approx m\left(x_{0}\right)+\epsilon_{i}$. Using this observation, to reduce the noise $\epsilon_{i}$, we can use the sample average. Thus, an estimator of $m\left(x_{0}\right)$ is to take the average of those responses whose covariate are close to $x_{0}$.

To make it more concrete, let $h>0$ be a threshold. The above procedure suggests to use

$$
\begin{equation*}
\widehat{m}_{\mathrm{loc}}\left(x_{0}\right)=\frac{\sum_{i:\left|X_{i}-x_{0}\right| \leq h} Y_{i}}{n_{h}\left(x_{0}\right)}=\frac{\sum_{i=1}^{n} Y_{i} I\left(\left|X_{i}-x_{0}\right| \leq h\right)}{\sum_{i=1}^{n} I\left(\left|X_{i}-x_{0}\right| \leq h\right)} \tag{3.1}
\end{equation*}
$$

where $n_{h}\left(x_{0}\right)$ is the number of observations where the covariate $X:\left|X_{i}-x_{0}\right| \leq h$. This estimator, $\widehat{m}_{\text {loc }}$, is called the local average estimator. Indeed, to estimate $m(x)$ at any given point $x$, we are using a local average as an estimator.

The local average estimator can be rewritten as

$$
\begin{equation*}
\widehat{m}_{\mathrm{loc}}\left(x_{0}\right)=\frac{\sum_{i=1}^{n} Y_{i} I\left(\left|X_{i}-x_{0}\right| \leq h\right)}{\sum_{i=1}^{n} I\left(\left|X_{i}-x_{0}\right| \leq h\right)}=\sum_{i=1}^{n} \frac{I\left(\left|X_{i}-x_{0}\right| \leq h\right)}{\sum_{\ell=1}^{n} I\left(\left|X_{\ell}-x_{0}\right| \leq h\right)} \cdot Y_{i}=\sum_{i=1}^{n} W_{i}\left(x_{0}\right) Y_{i} \tag{3.2}
\end{equation*}
$$

where

$$
\begin{equation*}
W_{i}\left(x_{0}\right)=\frac{I\left(\left|X_{i}-x_{0}\right| \leq h\right)}{\sum_{\ell=1}^{n} I\left(\left|X_{\ell}-x_{0}\right| \leq h\right)} \tag{3.3}
\end{equation*}
$$

is a weight for each observation. Note that $\sum_{i=1}^{n} W_{i}\left(x_{0}\right)=1$ and $W_{i}\left(x_{0}\right)>0$ for all $i=1, \cdots, n$; this implies that $W_{i}\left(x_{0}\right)$ 's are indeed weights. Equation (3.2) shows that the local average estimator can be written as a weighted average estimator so the $i$-th weight $W_{i}\left(x_{0}\right)$ determines the contribution of response $Y_{i}$ to the estimator $\widehat{m}_{\text {loc }}\left(x_{0}\right)$.

In constructing the local average estimator, we are placing a hard-thresholding on the neighboring pointsthose within a distance $h$ are given equal weight but those outside the threshold $h$ will be ignored completely. This hard-thresholding leads to an estimator that is not continuous.

To avoid problem, we consider another construction of the weights. Ideally, we want to give more weights to those observations that are close to $x_{0}$ and we want to have a weight that is 'smooth'. The Gaussian function $G(x)=\frac{1}{\sqrt{2 \pi}} e^{-x^{2} / 2}$ seems to be a good candidate. We now use the Gaussian function to construct an estimator. We first construct the weight

$$
W_{i}^{G}\left(x_{0}\right)=\frac{G\left(\frac{x_{0}-X_{i}}{h}\right)}{\sum_{\ell=1}^{n} G\left(\frac{x_{0}-X_{\ell}}{h}\right)} .
$$

The quantity $h>0$ is the similar quantity to the threshold in the local average but now it acts as the smoothing bandwidth of the Gaussian. After constructing the weight, our new estimator is

$$
\begin{equation*}
\widehat{m}_{\mathrm{G}}\left(x_{0}\right)=\sum_{i=1}^{n} W_{i}^{G}\left(x_{0}\right) Y_{i}=\sum_{i=1}^{n} \frac{G\left(\frac{x_{0}-X_{i}}{h}\right)}{\sum_{\ell=1}^{n} G\left(\frac{x_{0}-X_{\ell}}{h}\right)} Y_{i}=\frac{\sum_{i=1}^{n} Y_{i} G\left(\frac{x_{0}-X_{i}}{h}\right)}{\sum_{\ell=1}^{n} G\left(\frac{x_{0}-X_{\ell}}{h}\right)} . \tag{3.4}
\end{equation*}
$$

This new estimator has a weight that changes more smoothly than the local average and is smooth as we desire.

Observing from equation (3.1) and (3.4), one may notice that these local estimators are all of a similar form:

$$
\begin{equation*}
\widehat{m}_{h}\left(x_{0}\right)=\frac{\sum_{i=1}^{n} Y_{i} K\left(\frac{x_{0}-X_{i}}{h}\right)}{\sum_{\ell=1}^{n} K\left(\frac{x_{0}-X_{\ell}}{h}\right)}=\sum_{i=1}^{n} W_{i}^{K}\left(x_{0}\right) Y_{i}, \quad W_{i}^{K}\left(x_{0}\right)=\frac{K\left(\frac{x_{0}-X_{i}}{h}\right)}{\sum_{\ell=1}^{n} K\left(\frac{x_{0}-X_{\ell}}{h}\right)}, \tag{3.5}
\end{equation*}
$$

where $K$ is some function. When $K$ is a Gaussian, we obtain estimator (3.4); when $K$ is a uniform over $[-1,1]$, we obtain the local average (3.1). The estimator in equation (3.5) is called the kernel regression estimator or Nadaraya-Watson estimator ${ }^{1}$. The function $K$ plays a similar role as the kernel function in the KDE and thus it is also called the kernel function. And the quantity $h>0$ is similar to the smoothing bandwidth in the KDE so it is also called the smoothing bandwidth.

### 3.3.1 Theory

Now we study some statistical properties of the estimator $\widehat{m}_{h}$. Suppose that we are interested in $m(x)$ over a compact interval $\mathbb{K} \subset \mathbb{R}$.

Theorem 3.2 Assume that

- $\inf _{x \in \mathbb{K}} p(x)>0$ and $p(x)$ has bounded second derivatives.
- $\mathbb{E}\left(Y^{2} \mid X=x\right)<\infty$ and $m(x)$ has bounded third derivatives.

Then

$$
\begin{aligned}
\operatorname{bias}\left(\widehat{m}_{h}(x)\right) & =\frac{h^{2}}{2} \mu_{K}\left(m^{\prime \prime}(x)+2 \frac{m^{\prime}(x) p^{\prime}(x)}{p(x)}\right)+o\left(h^{2}\right) \\
\operatorname{Var}\left(\widehat{m}_{h}(x)\right) & =\frac{\sigma^{2} \cdot \sigma_{K}^{2}}{p(x)} \cdot \frac{1}{n h}+o\left(\frac{1}{n h}\right),
\end{aligned}
$$

where $\mu_{K}=\int x^{2} K(x) d x$ is the same constant of the kernel function as in the $K D E$ and $\sigma^{2}=\operatorname{Var}\left(\epsilon_{i}\right)$ is the error of the regression model and $\sigma_{K}^{2}=\int K^{2}(x) d x$ is a constant of the kernel function (the same as in the KDE).

[^0]The bias has two components: a curvature component $m^{\prime \prime}(x)$ and a design component $\frac{m^{\prime}(x) p^{\prime}(x)}{p(x)}$. The curvature component is similar to the one in the KDE; when the regression function curved a lot, kernel smoothing will smooth out the structure, introducing some bias. The second component, also known as the design bias, is a new component compare to the bias in the KDE. This component depends on the density of covariate $p(x)$. Note that in some studies, we can choose the values of covariates so the density $p(x)$ is also called the design (this is why it is known as the design bias).

The expression of variance tells us possible sources of variability. First, the variance increases when $\sigma^{2}$ increases. This makes perfect sense because $\sigma^{2}$ is the noise level. When the noise level is large, we expect the estimation error increases. Second, the density of covariate $p(x)$ is inversely related to the variance. This is also very reasonable because when $p(x)$ is large, there tends to be more data points around $x$, increasing the size of sample that we are averaging from. Last, the convergence rate is $O\left(\frac{1}{n h}\right)$, which is the same as the KDE.

MSE and MISE. Using the expression of bias and variance, the MSE at point $x$ is

$$
\operatorname{MSE}\left(\widehat{m}_{h}(x)\right)=\frac{h^{4}}{4} \mu_{K}^{2}\left(m^{\prime \prime}(x)+2 \frac{m^{\prime}(x) p^{\prime}(x)}{p(x)}\right)^{2}+\frac{\sigma^{2} \cdot \sigma_{K}^{2}}{p(x)} \cdot \frac{1}{n h}+o\left(h^{4}\right)+o\left(\frac{1}{n h}\right)
$$

and the MISE is

$$
\begin{equation*}
\operatorname{MISE}\left(\widehat{m}_{h}\right)=\frac{h^{4}}{4} \mu_{K}^{2} \int\left(m^{\prime \prime}(x)+2 \frac{m^{\prime}(x) p^{\prime}(x)}{p(x)}\right)^{2} d x+\frac{\sigma^{2} \cdot \sigma_{K}^{2}}{n h} \int \frac{1}{p(x)} d x+o\left(h^{4}\right)+o\left(\frac{1}{n h}\right) . \tag{3.6}
\end{equation*}
$$

Optimizing the major components in equation (3.6) (the AMISE), we obtain the optimal value of the smoothing bandwidth

$$
h_{\mathrm{opt}}=C^{*} \cdot n^{-1 / 5},
$$

where $C^{*}$ is a constant depending on $p$ and $K$.

### 3.3.2 Cross-Validation

The smoothing bandwidth $h$ has to be chosen to construct our estimator. The theory suggests that we choose it to be $h_{\mathrm{opt}}=C^{*} \cdot n^{-1 / 5}$ but this involves an unknown quantity $C^{*}$. So in practice, how can we choose it? A good news is that-unlike the density estimation problem, there is a simple approach to choose $h$ : the cross-validation $(\mathrm{CV})^{2}$.

Before we discuss the details of CV, we first introduce the predictive risk. Let $\widehat{m}_{h}$ be the kernel regression using the $n$ observations. Let $X_{n+1}, Y_{n+1}$ be a new observation (from the same population). We define the predictive risk of our regression estimator as

$$
\begin{equation*}
R(h)=\mathbb{E}\left(Y_{n+1}-\widehat{m}_{h}\left(X_{n+1}\right)\right)^{2} . \tag{3.7}
\end{equation*}
$$

Namely, the quantity $R(h)$ is the expected square error of predicting the next observation using the kernel regression.

CV is a collection of approaches that tries to estimate the predictive risk $R(h)$ using a data-splitting approach. A classical version of CV is the leave-one out cross-validation (LOO-CV):

$$
\widehat{R}(h)=\frac{1}{n} \sum_{i=1}^{n}\left(Y_{i}-\widehat{m}_{h,-i}\left(X_{i}\right)\right)^{2}
$$

[^1]where $\widehat{m}_{h,-i}\left(X_{i}\right)$ is the kernel regression using all observations except $i$-th observation $X_{i}, Y_{i}$. Namely, LOOCV leaves each observation out once at a time and use the remaining observations to train the estimator and evaluate the quality of the estimator using the left out observation. The main reasoning of such a procedure is to make sure we do not use the data twice.

Another popular version of CV is the K-fold CV. We randomly split the data into K equal size groups. Each time we leave out one group and use the other K-1 groups to construct our estimator. Then we use the left out group to evaluate the risk. Repeat this procedure many times and take the average as the risk estimator $\widehat{R}(h)$.

## K-Fold Cross-Validation.

1. Randomly split $\mathcal{D}=\left\{\left(X_{1}, Y_{1}\right), \cdots,\left(X_{n}, Y_{n}\right)\right\}$ into K groups: $\mathcal{D}_{1}, \cdots, \mathcal{D}_{K}$.
2. For $\ell$-th group, construct the estimator $\widehat{m}_{h}^{(\ell)}$ using all the data except $\ell$-th group.
3. Evaluate the error by

$$
\widehat{R}^{(\ell)}(h)=\frac{1}{n_{\ell}} \sum_{\left(X_{i}, Y_{i}\right) \in \mathcal{D}_{\ell}}\left(Y_{i}-\widehat{m}_{h}^{(\ell)}\left(X_{i}\right)\right)^{2}
$$

4. Compute the average error

$$
\widehat{R}(h)=\frac{1}{K} \sum_{\ell=1}^{K} \widehat{R}^{(\ell)}(h) .
$$

5. Repeat the above 4 steps $N$ times, leading to N average errors

$$
\widehat{R}^{*(1)}(h), \cdots, \widehat{R}^{*(N)}(h) .
$$

6. Estimate $R(h)$ via

$$
\widehat{R}^{*}(h)=\frac{1}{N} \sum_{\ell=1}^{N} \widehat{R}^{*(\ell)}(h) .
$$

The CV provides a simple approach of estimating the predictive errors. To choose the smoothing bandwidth, we pick

$$
h^{*}=\operatorname{argmin}_{h>0} \widehat{R}(h) .
$$

In practice, we apply the CV to various values of $h$ and choose the one with the minimal predictive risk. Generally, we will plot $\widehat{R}(h)$ versus $h$ and determine if the minimal value makes sense. Sometimes there might be no well-defined minimum value (like a flat region).

Why do we want to split the data into two parts and construct the estimator on one part and evaluate the risk on the other part? The main reason is to obtain a reliable estimate of the predictive risk. If we use the same set of data to construct our estimator and evaluate the errors, the estimated predictive risk will be smaller than the actual predictive risk. To see this, consider the local average estimator with $h \approx 0$. When $h$ is very very small, $\widehat{m}_{\text {loc }}\left(X_{i}\right)=Y_{i}$ because the neighborhood only contain this single observation. In this case, the estimated predictive risk will be $\sum_{i=1}^{n}\left(Y_{i}-Y_{i}\right)^{2}=0$. This is related to the so-called overfitting ${ }^{3}$.

The cross-validation is a very popular and common approach to choose a tuning parameter. Other tuning

[^2]parameters such as the number of basis $N$ and the penalization $\lambda$ (will be introduced in a minute) can all be chosen by minimizing the cross-validation error.

### 3.3.3 Uncertainty and Confidence Intervals

How do we assess the quality of our estimator $\widehat{m}_{h}(x)$ ?
We can use the bootstrap to do it. In this case, empirical bootstrap, residual bootstrap, and wild bootstrap all can be applied. But note that each of them relies on slightly different assumptions. Let $\left(X_{1}^{*}, Y_{1}^{*}\right), \cdots,\left(X_{n}^{*}, Y_{n}^{*}\right)$ be the bootstrap sample. Applying the bootstrap sample to equation (3.5), we obtain a bootstrap kernel regression, denoted as $\widehat{m}_{h}^{*}$. Now repeat the bootstrap procedure $B$ times, this yields

$$
\widehat{m}_{h}^{*(1)}, \cdots, \widehat{m}_{h}^{*(B)},
$$

$B$ bootstrap kernel regression estimator. Then we can estimate the variance of $\widehat{m}_{h}(x)$ by the sample variance

$$
\widehat{\operatorname{Var}}_{B}\left(\widehat{m}_{h}(x)\right)=\frac{1}{B-1} \sum_{\ell=1}^{B}\left(\widehat{m}_{h}^{*(\ell)}(x)-\overline{\widehat{m}}_{h, B}^{*}(x)\right), \quad \overline{\widehat{m}}_{h, B}^{*}(x)=\frac{1}{B} \sum_{\ell=1}^{B} \widehat{m}_{h}^{*(\ell)}(x)
$$

Similarly, we can estimate the MSE as what we did in Lecture 5 and 6. However, when using the bootstrap to estimate the uncertainty, one has to be very careful because when $h$ is either too small or too large, the bootstrap estimate may fail to converge its target.

When we choose $h=O\left(n^{-1 / 5}\right)$, the bootstrap estimate of the variance is consistent but the bootstrap estimate of the MSE might not be consistent. The main reason is: it is easier for the bootstrap to estimate the variance than the bias. Thus, when we choose $h$ in such a way, both bias and the variance contribute a lot to the MSE so we cannot ignore the bias. However, in this case, the bootstrap cannot estimate the bias consistently so the estimate of the MSE is not consistent.

Confidence interval. To construct a confidence interval of $m(x)$, we will use the following property of the kernel regression.

Theorem 3.3 Assume the conditions in Theorem 3.2. Then

$$
\begin{aligned}
\sqrt{n h}\left(\widehat{m}_{h}(x)-\mathbb{E}\left(\widehat{m}_{h}(x)\right)\right) & \xrightarrow{D} N\left(0, \frac{\sigma^{2} \cdot \sigma_{K}^{2}}{p(x)}\right) \\
\frac{\widehat{m}_{h}(x)-\mathbb{E}\left(\widehat{m}_{h}(x)\right.}{\operatorname{Var}\left(\widehat{m}_{h}(x)\right)} & \xrightarrow{D} N(0,1) .
\end{aligned}
$$

The variance depends on three quantities: $\sigma^{2}, \sigma_{K}^{2}$, and $p(x)$. The quantity $\sigma_{K}^{2}$ is known because it is just a characteristic of the kernel function. The density of covariates $p(x)$ can be estimated using a KDE. So what remains unknown is the noise level $\sigma^{2}$. A good news is: we can estimate it using the residuals. Recall that residuals are

$$
e_{i}=Y_{i}-\widehat{Y}_{i}=Y_{i}-\widehat{m}_{h}\left(X_{i}\right)
$$

When $\widehat{m}_{h} \approx m$, the residual becomes an approximation to the noise $\epsilon_{i}$. The quantity $\sigma^{2}=\operatorname{Var}\left(\epsilon_{1}\right)$ so we can use the sample variance of the residuals to estimate it (note that the average of residuals is 0 ):

$$
\begin{equation*}
\widehat{\sigma}^{2}=\frac{1}{n-2 \nu+\widetilde{\nu}} \sum_{i=1}^{n} e_{i}^{2} \tag{3.8}
\end{equation*}
$$

where $\nu, \widetilde{\nu}$ are quantities acting as degree-of-freedom in which we will explain later. Thus, a $1-\alpha$ CI can be constructed using

$$
\widehat{m}_{h}(x) \pm z_{1-\alpha / 2} \frac{\widehat{\sigma} \cdot \sigma_{K}}{\sqrt{\widehat{p}_{n}(x)}}
$$

where $\widehat{p}_{n}(x)$ is the KDE of the covariates.
However, note that this confidence interval is a pointwise confidence interval, i.e., we are controlling the type-1 error at a given point $x$. Also, this confidence interval has a limitation that the bias may decrease our coverage- the construction shows that our confidence interval has the right coverage for $\mathbb{E}\left(\widehat{m}_{h}(x)\right)$, not the true regression function $m(x)$. Thus, we often need to undersmooth the data a bit to obtain a proper coverage. One possible solution to this problem is to use the debiased estimator; see the following paper for more information

Cheng, G., \& Chen, Y. C. (2019). Nonparametric inference via bootstrapping the debiased estimator. Electronic Journal of Statistics, 13(1), 2194-2256.

### 3.3.4 Relation to KDE

Many theoretical results of the KDE apply to the nonparametric regression. For instance, we can generalize the MISE into other types of error measurement between $\widehat{m}_{h}$ and $m$. We can also use derivatives of $\widehat{m}_{h}$ as estimators of the corresponding derivatives of $m$. Moreover, when we have a multivariate covariate, we can use either a radial basis kernel or a product kernel to generalize the kernel regression to multivariate case.

The KDE and the kernel regression has a very interesting relationship. Using the given bivariate random sample $\left(X_{1}, Y_{1}\right), \cdots,\left(X_{n}, Y_{n}\right)$, we can estimate the joint PDF $p(x, y)$ as

$$
\widehat{p}_{n}(x, y)=\frac{1}{n h^{2}} \sum_{i=1}^{n} K\left(\frac{X_{i}-x}{h}\right) K\left(\frac{Y_{i}-y}{h}\right) .
$$

This joint density estimator also leads to a marginal density estimator of $X$ :

$$
\widehat{p}_{n}(x)=\int \widehat{p}_{n}(x, y) d y=\frac{1}{n h} \sum_{i=1}^{n} K\left(\frac{X_{i}-x}{h}\right) .
$$

Theorem 3.4 The conditional expectation of $Y$ given $X=x$ implied by the $2 D K D E$ is the same as the kernel regression estimator.

Proof: Recalled that the regression function is the conditional expectation

$$
m(x)=\mathbb{E}(Y \mid X=x)=\int y p(y \mid x) d y=\int y \frac{p(x, y)}{p(x)} d y=\frac{\int y p(x, y) d y}{p(x)}
$$

Replacing $p(x, y)$ and $p(x)$ by their corresponding estimators $\widehat{p}_{n}(x, y)$ and $\widehat{p}_{n}(x)$, we obtain an estimate of
$m(x)$ as

$$
\begin{aligned}
\widehat{m}_{n}(x) & =\frac{\int y \widehat{p}_{n}(x, y) d y}{\widehat{p}_{n}(x)} \\
& =\frac{\int y \frac{1}{n h^{2}} \sum_{i=1}^{n} K\left(\frac{X_{i}-x}{h}\right) K\left(\frac{Y_{i}-y}{h}\right) d y}{\frac{1}{n h} \sum_{i=1}^{n} K\left(\frac{X_{i}-x}{h}\right)} \\
& =\frac{\sum_{i=1}^{n} K\left(\frac{X_{i}-x}{h}\right) \cdot \int y \cdot K\left(\frac{Y_{i}-y}{h}\right) \frac{d y}{h}}{\sum_{i=1}^{n} K\left(\frac{X_{i}-x}{h}\right)} \\
& =\frac{\sum_{i=1}^{n} K\left(\frac{X_{i}-x}{h}\right) Y_{i}}{\sum_{i=1}^{n} K\left(\frac{X_{i}-x}{h}\right)} \\
& =\frac{\sum_{i=1}^{n} Y_{i} K\left(\frac{X_{i}-x}{h}\right)}{\sum_{i=1}^{n} K\left(\frac{X_{i}-x}{h}\right)} \\
& =\widehat{m}_{h}(x)
\end{aligned}
$$

Note that when $K(x)$ is symmetric, $\int y \cdot K\left(\frac{Y_{i}-y}{h}\right) \frac{d y}{h}=Y_{i}$.
Namely, we may understand the kernel regression as an estimator inverting the KDE of the joint PDF into a regression estimator.

### 3.4 Local Polynomial Regression

The kernel regression estimator has a limitation that it suffers a lot from the boundary bias, i.e., when $x$ is close to the support of $p(x)$, the bias will be very large. To address this issue, we may use a modified estimator called the local polynomial regression (LPR).

LPR starts with the following localized least squared estimation problem. Suppose that we want to estimate the regression function $m(x)$ at point $x$. Consider fitting the following local linear function

$$
\operatorname{LPR}\left(\beta_{0}, \beta_{1} ; x\right)=\sum_{i=1}^{n} K\left(\frac{x-X_{i}}{h}\right)\left(Y_{i}-\beta_{0}-\beta_{1} X_{i}\right)^{2}
$$

where $K(\cdot)$ is the usual kernel function. This is fitting a weighted linear regression where the observations close to $x$ are given higher weights. Let $\widehat{\beta}_{0}(x), \widehat{\beta}_{1}(x)$ be the minimizer of $\operatorname{LPR}\left(\beta_{0}, \beta_{1} ; x\right)$. Then the estimator $\widehat{\beta}_{0}(x)$ is called the local linear smoother and is a consistent estimator of $m(x)$, the regression function.

There is a closed-form of the local linear smoother. Define the diagonal matrix

$$
W(x) \in \mathbb{R}^{n \times n}=\operatorname{Diag}\left(K\left(\frac{x-X_{1}}{h}\right), \cdots, K\left(\frac{x-X_{n}}{h}\right)\right)
$$

and the matrix $\mathbb{X} \in \mathbb{R}^{n \times 2}$

$$
\mathbb{X}=\left(1_{n}, X\right)
$$

where $X$ is a column vector of $X_{1}, \cdots, X_{n}$ and $1_{n}$ is a column vector of 1 's. Let $\mathbb{Y}$ be the column vector of $Y_{1}, \cdots, Y_{n}$. Using the derivation as linear regression, you can show that the local linear smoother $\widehat{\beta}_{0}(x)$ is

$$
\widehat{\beta}_{0}(x)=e_{1}^{T}\left(\mathbb{X}^{T} W(x) \mathbb{X}\right)^{-1} \mathbb{X}^{T} W(x) \mathbb{Y}
$$

where $e_{1}^{T}=(1,0)$ is a $1 \times 2$ vector. We write $\widehat{m}_{\mathrm{LL}}(x)=\widehat{\beta}_{0}(x)$.
We can generalize the local linear smoother to higher order polynomials. For instance, we can fit it to the $q$-th order polynomial

$$
\operatorname{LPR}\left(\beta_{0}, \beta_{1}, \cdots, \beta_{q} ; x\right)=\sum_{i=1}^{n} K\left(\frac{x-X_{i}}{h}\right)\left(Y_{i}-\beta_{0}-\beta_{1} X_{i}-\cdots \beta_{q} X_{i}^{q}\right)^{2}
$$

Fitting a higher order polynomial is often used to estimate the derivative of the regression function. In fact, if we want to consistently estimate $m^{(\beta)}$, the $\beta$-the derivative of $m(x)$, then we will fit a $(\beta+1)$-th order polynomial. Using the same derivation as the local linear smoother, you can obtain a closed-form of the estimator.

Theorem 3.5 (Fan (1992)) Suppose that $Y_{i}=m\left(X_{i}\right)+\sigma\left(X_{i}\right) \epsilon_{i}$ and $X_{i} \in \mathbb{K}$ with $\mathbb{E}\left(\epsilon_{i}\right)=0, \operatorname{Var}\left(\epsilon_{i}\right)=1$ and $X_{1}, \cdots, X_{n} \sim p$ and we are interested in a point $x$ in the interior of $\mathbb{K}$. Assume the followings:

- $p(x)>0$.
- $p, m^{\prime \prime}$, and $\sigma$ are continuous in the neighborhood of $x$.
- $h \rightarrow 0, n h \rightarrow \infty$.

Then the local linear smoother satisfies

$$
\operatorname{bias}\left(\widehat{m}_{\mathrm{LL}}(x)\right)=\frac{h^{2}}{2} m^{\prime \prime}(x) \mu_{K}+o\left(h^{2}\right), \quad \operatorname{Var}\left(\widehat{m}_{\mathrm{LL}}(x)\right)=\frac{\sigma^{2}(x)}{p(x) n h} \sigma_{K}^{2}+o\left(\frac{1}{n h}\right) .
$$

Namely, the local linear smoother does not suffer from the design bias. The above theorem is from the following paper:

Fan, J. (1992). Design-adaptive nonparametric regression. Journal of the American statistical Association, 87(420), 998-1004.

### 3.5 Linear Smoother

Now we are going to introduce a very important notion called linear smoother. Linear smoother is a collection of many regression estimators that have nice properties. The linear smoother is an estimator of the regression function in the form that

$$
\begin{equation*}
\widehat{m}(x)=\sum_{i=1}^{n} \ell_{i}(x) Y_{i}, \tag{3.9}
\end{equation*}
$$

where $\ell_{i}(x)$ is some function depending on $X_{1}, \cdots, X_{n}$ but not on any of $Y_{1}, \cdots, Y_{n}$.
The residual for the $j$-th observation can be written as

$$
e_{j}=Y_{j}-\widehat{m}\left(X_{j}\right)=Y_{j}-\sum_{i=1}^{n} \ell_{i}\left(X_{j}\right) Y_{i}
$$

Let $e=\left(e_{1}, \cdots, e_{n}\right)^{T}$ be the vector of residuals and define an $n \times n$ matrix $L$ as $L_{i j}=\ell_{j}\left(X_{i}\right)$ :

$$
L=\left(\begin{array}{ccccc}
\ell_{1}\left(X_{1}\right) & \ell_{2}\left(X_{1}\right) & \ell_{3}\left(X_{1}\right) & \cdots & \ell_{n}\left(X_{1}\right) \\
\ell_{1}\left(X_{2}\right) & \ell_{2}\left(X_{2}\right) & \ell_{3}\left(X_{2}\right) & \cdots & \ell_{n}\left(X_{2}\right) \\
\vdots & \vdots & \vdots & \vdots & \vdots \\
\ell_{1}\left(X_{n}\right) & \ell_{2}\left(X_{n}\right) & \ell_{3}\left(X_{n}\right) & \cdots & \ell_{n}\left(X_{n}\right)
\end{array}\right)
$$

Then the predicted vector $\widehat{\mathbb{Y}}=\left(\widehat{Y}_{1}, \cdots, \widehat{Y}_{n}\right)^{T}=L Y$, where $\mathbb{Y}=\left(Y_{1}, \cdots, Y_{n}\right)^{T}$ is the vector of observed $Y_{i}$ 's and $e=\mathbb{Y}-\widehat{\mathbb{Y}}=\mathbb{Y}-L \mathbb{Y}=(I-L) \mathbb{Y}$.

Example: Linear Regression. For the linear regression, let $\mathbb{X}$ denotes the data matrix (first column is all value 1 and second column is $\left.X_{1}, \cdots, X_{n}\right)$. We know that $\widehat{\beta}=\left(\mathbb{X}^{T} \mathbb{X}\right)^{-1} \mathbb{X}^{T} \mathbb{Y}$ and $\widehat{\mathbb{Y}}=\mathbb{X} \widehat{\beta}=\mathbb{X}\left(\mathbb{X}^{T} \mathbb{X}\right)^{-1} \mathbb{X}^{T} \mathbb{Y}$. This implies that the matrix $L$ is

$$
L=\mathbb{X}\left(\mathbb{X}^{T} \mathbb{X}\right)^{-1} \mathbb{X}^{T}
$$

which is also the projection matrix in linear regression. Thus, the linear regression is a linear smoother.
Example: Regressogram. The regressogram is also a linear smoother. Let $B_{1}, \cdots, B_{m}$ be the bins of the covariate and define $B(x)$ be the bin such that $x$ belongs to. Then

$$
\ell_{j}(x)=\frac{I\left(X_{j} \in B(x)\right)}{\sum_{i=1}^{n} I\left(X_{i} \in B(x)\right)}
$$

Example: Kernel Regression. As you may expect, the kernel regression is also a linear smoother. Recall from equation (3.5)

$$
\widehat{m}_{h}\left(x_{0}\right)=\frac{\sum_{i=1}^{n} Y_{i} K\left(\frac{x_{0}-X_{i}}{h}\right)}{\sum_{\ell=1}^{n} K\left(\frac{x_{0}-X_{\ell}}{h}\right)}=\sum_{i=1}^{n} W_{i}^{K}\left(x_{0}\right) Y_{i}, \quad W_{i}^{K}\left(x_{0}\right)=\frac{K\left(\frac{x_{0}-X_{i}}{h}\right)}{\sum_{\ell=1}^{n} K\left(\frac{x_{0}-X_{\ell}}{h}\right)}
$$

so

$$
\ell_{j}(x)=\frac{K\left(\frac{x-X_{j}}{h}\right)}{\sum_{\ell=1}^{n} K\left(\frac{x-X_{\ell}}{h}\right)} .
$$

Example: Local Linear Smoother. You can easily show that the LPR is a linear smoother. In particular, the linear smoother has the vector $\ell(x)=\left(\ell_{1}(x), \cdots, \ell_{n}(x)\right)^{T}$ as

$$
\ell(x)=e_{1}^{T}\left(\mathbb{X}^{T} W(x) \mathbb{X}\right)^{-1} \mathbb{X}^{T} W(x)
$$

### 3.5.1 Variance of Linear Smoother

The linear smoother has an unbiased estimator of the underlying noise level $\sigma^{2}$ under the fixed design, i.e., the covariates are non-random. Recall that then noise level $\sigma^{2}=\operatorname{Var}\left(\epsilon_{i}\right)$.

We need to use two tricks about variance and covariance matrix. For a matrix $A$ and a random variable $X$,

$$
\operatorname{Cov}(A X)=A \operatorname{Cov}(X) A^{T} .
$$

Thus, the covariance matrix of the residual vector

$$
\operatorname{Cov}(e)=\operatorname{Cov}((I-L) \mathbb{Y})=(I-L) \operatorname{Cov}(\mathbb{Y})\left(I-L^{T}\right)
$$

Because $Y_{1}, \cdots, Y_{n}$ are IID, $\operatorname{Cov}(\mathbb{Y})=\sigma^{2} \mathbb{I}_{n}$, where $\mathbb{I}_{n}$ is the $n \times n$ identity matrix. This implies

$$
\operatorname{Cov}(e)=(I-L) \operatorname{Cov}(\mathbb{Y})\left(I-L^{T}\right)=\sigma^{2}\left(I-L-L^{T}+L L^{T}\right)
$$

Now taking matrix trace in both side,

$$
\operatorname{Tr}(\operatorname{Cov}(e))=\sum_{i=1}^{n} \operatorname{Var}\left(e_{i}\right)=\sigma^{2} \operatorname{Tr}\left(I-L-L^{T}+L L^{T}\right)=\sigma^{2}(n-\nu-\nu+\widetilde{\nu})
$$

where $\nu=\operatorname{Tr}(L)$ and $\widetilde{\nu}=\operatorname{Tr}\left(L L^{T}\right)$. Because the residual square is approximately $\operatorname{Var}\left(e_{i}\right)$, we have

$$
\sum_{i=1}^{n} e_{i}^{2} \approx \sum_{i=1}^{n} \operatorname{Var}\left(e_{i}\right)=\sigma^{2}(n-2 \nu+\widetilde{\nu})
$$

Thus, we can estimate $\sigma^{2}$ by

$$
\begin{equation*}
\widehat{\sigma}^{2}=\frac{1}{n-2 \nu+\widetilde{\nu}} \sum_{i=1}^{n} e_{i}^{2} \tag{3.10}
\end{equation*}
$$

which is what we did in equation (3.8). The quantity $\nu$ is called the degree of freedom. In the linear regression case, $\nu=\widetilde{\nu}=p+1$, the number of covariates so the variance estimator $\widehat{\sigma}^{2}=\frac{1}{n-p-1} \sum_{i=1}^{n} e_{i}^{2}$. If you have learned the variance estimator of a linear regression, you should be familiar with this estimator.

The degree of freedom $\nu$ is easy to interpret in the linear regression. And the power of equation (3.10) is that it works for every linear smoother as long as the errors $\epsilon_{i}$ 's are IID. So it shows how we can define effective degree of freedom for other complicated regression estimator.

### 3.6 Basis Approach

Recall that we observes pairs $\left(X_{1}, Y_{1}\right), \cdots,\left(X_{n}, Y_{n}\right)$ and we are interested in the regression function $m(x)=$ $\mathbb{E}\left(Y_{1} \mid X_{1}=x\right)$. In this section, we will make the following two assumptions:

- $Y_{i}=m\left(X_{i}\right)+\sigma \cdot \epsilon_{i}$, where $\epsilon_{i} \sim N(0,1)$ is the noise. Moreover, $\epsilon_{1}, \cdots, \epsilon_{n}$ are IID.
- $X_{i}=\frac{i}{n}$. Namely, the covariates consist a uniform grid over $[0,1]$ and is non-random.

Similar to the basis approach for the density estimation problem where we approximate the density function by the sum of coefficients and basis, we will approximate the regression function by a basis:

$$
m(x)=\sum_{j=1}^{\infty} \theta_{j} \phi_{j}(x)
$$

where $\left\{\phi_{1}, \phi_{2}, \cdots\right\}$ is an orthonormal basis and $\theta_{1}, \theta_{2}, \cdots$ are the coefficients.
Again, here we consider the cosine basis:

$$
\phi_{1}(x)=1, \quad \phi_{j}(x)=\sqrt{2} \cos ((j-1) \pi x), j=2,3, \cdots
$$

As is done in the density estimation, we will use only the top $M$ basis to form our estimator. Namely,

$$
\widehat{m}_{M}(x)=\sum_{j=1}^{M} \widehat{\theta}_{j} \phi_{j}(x),
$$

for some coefficient estimates $\widehat{\theta}_{1}, \cdots$. Again, $M$ is the tuning parameter in our estimator.
Here is a simple choice of the coefficient estimates that we will be using:

$$
\widehat{\theta}_{j}=\frac{1}{n} \sum_{i=1}^{n} Y_{i} \phi_{j}\left(X_{i}\right)=\frac{1}{n} \sum_{i=1}^{n} Y_{i} \phi_{j}\left(\frac{i}{n}\right) .
$$

To determine the tuning parameter $M$, we analyze the MISE. We start with analyzing the bias and variance of $\widehat{\theta}_{j}$.

### 3.6.1 Asymptotic theory

Asymptotic normality. Note that the estimator can be rewritten as

$$
\begin{aligned}
\widehat{m}_{M}(x) & =\sum_{j=1}^{M} \widehat{\theta}_{j} \phi_{j}(x) \\
& =\sum_{j=1}^{M} \frac{1}{n} \sum_{i=1}^{n} Y_{i} \phi_{j}\left(\frac{i}{n}\right) \phi_{j}(x) \\
& =\frac{1}{n} \sum_{i=1}^{n} Y_{i} \sum_{j=1}^{M} \phi_{j}\left(\frac{i}{n}\right) \phi_{j}(x) .
\end{aligned}
$$

Thus, for $M$ being fixed, we have

$$
\sqrt{n}\left(\widehat{m}_{M}(x)-\mathbb{E}\left(\widehat{m}_{M}(x)\right)\right) \xrightarrow{D} N\left(0, \sigma_{M}^{2}\right)
$$

for some $\sigma_{M}^{2}$. Note that later our analysis will demonstrate

$$
\mathbb{E}\left(\widehat{m}_{M}(x)\right)=\sum_{j=1}^{M} \theta_{j} \phi_{j}(x), \quad \sigma_{M}^{2}=\sigma^{2} \sum_{j=1}^{M} \phi_{j}^{2}(x)
$$

## Bias.

$$
\begin{aligned}
\operatorname{bias}\left(\widehat{\theta}_{j}\right) & =\mathbb{E}\left(\widehat{\theta}_{j}\right)-\theta_{j} \\
& =\mathbb{E}\left(\left.\frac{1}{n} \sum_{i=1}^{n} Y_{i} \phi_{j}\left(\frac{i}{n}\right) \right\rvert\, X_{i}=\frac{i}{n}\right)-\theta_{j} \\
& =\frac{1}{n} \sum_{i=1}^{n} \mathbb{E}\left(Y_{i} \left\lvert\, X_{i}=\frac{i}{n}\right.\right) \phi_{j}\left(\frac{i}{n}\right)-\theta_{j} \\
& =\frac{1}{n} \sum_{i=1}^{n} m\left(\frac{i}{n}\right) \phi_{j}\left(\frac{i}{n}\right)-\theta_{j} \\
& =\frac{1}{n} \sum_{i=1}^{n} m\left(\frac{i}{n}\right) \phi_{j}\left(\frac{i}{n}\right)-\int_{0}^{1} m(x) \phi_{j}(x) d x
\end{aligned}
$$

Namely, the bias is the difference between actual integration and a discretized version of integration. We know that when $n$ is large, the two integrations are almost the same so we can ignore the bias. Thus, we will write

$$
\operatorname{bias}\left(\widehat{\theta}_{j}\right)=0
$$

for simplicity.

## Variance.

$$
\begin{aligned}
\operatorname{Var}\left(\widehat{\theta}_{j}\right) & =\operatorname{Var}(\frac{1}{n} \sum_{i=1}^{n} \underbrace{\left(m\left(\frac{i}{n}\right)+\sigma \cdot \epsilon_{i}\right)}_{=Y_{i}} \phi_{j}\left(\frac{i}{n}\right)) \\
& =\frac{1}{n^{2}} \sum_{i=1}^{n} \phi_{j}^{2}\left(\frac{i}{n}\right) \operatorname{Var}\left(\epsilon_{i}\right) \sigma^{2} \\
& =\frac{\sigma^{2}}{n^{2}} \sum_{i=1}^{n} \phi_{j}^{2}\left(\frac{i}{n}\right) .
\end{aligned}
$$

Note that $\frac{1}{n} \sum_{i=1}^{n} \phi_{j}^{2}\left(\frac{i}{n}\right) \approx \int_{0}^{1} \phi_{j}^{2}(x) d x=1$. For simplicity, we just write

$$
\operatorname{Var}\left(\widehat{\theta}_{j}\right)=\frac{\sigma^{2}}{n}
$$

MISE. To analyze the MISE, we first note that the bias of $\widehat{m}_{M}(x)$ is

$$
\operatorname{bias}\left(\widehat{m}_{M}(x)\right)=\mathbb{E}\left(\widehat{m}_{M}(x)\right)-m(x)=\sum_{j=1}^{M} \theta_{j} \phi_{j}(x)-\sum_{j=1}^{\infty} \theta_{j} \phi_{j}(x)=\sum_{j=M+1}^{\infty} \theta_{j} \phi_{j}(x)
$$

This further implies that the integrated sqaured bias

$$
\begin{aligned}
\int_{0}^{1} \operatorname{bias}^{2}\left(\widehat{m}_{M}(x)\right) d x & =\int_{0}^{1} \sum_{j=M+1}^{\infty} \theta_{j} \phi_{j}(x) \sum_{\ell=M+1}^{\infty} \theta_{\ell} \phi_{\ell}(x) d x \\
& =\sum_{j=M+1}^{\infty} \theta_{j} \sum_{\ell=M+1}^{\infty} \theta_{\ell} \underbrace{\int_{0}^{1} \phi_{j}(x) \phi_{\ell}(x) d x}_{=I(j=\ell)} \\
& =\sum_{j=M+1}^{\infty} \theta_{j}^{2}
\end{aligned}
$$

Again, if we assume that $m$ satisfies $\int_{0}^{1}\left|m^{\prime \prime}(x)\right|^{2} d x<\infty$, we have

$$
\sum_{j=M+1}^{\infty} \theta_{j}^{2}=O\left(M^{-4}\right)
$$

Now we turn to the analysis of variance.

$$
\begin{aligned}
\operatorname{Var}\left(\widehat{m}_{M}(x)\right) & =\operatorname{Var}\left(\sum_{j=1}^{M} \widehat{\theta}_{j} \phi_{j}(x)\right) \\
& =\sum_{j=1}^{M} \operatorname{Var}\left(\widehat{\theta}_{j}\right) \phi_{j}^{2}(x)+\sum_{j \neq k}^{M} \operatorname{Cov}\left(\widehat{\theta}_{j}, \widehat{\theta}_{k}\right) \phi_{j}(x) \phi_{k}(x) \\
& =\frac{\sigma^{2}}{n} \sum_{j=1}^{M} \phi_{j}^{2}(x)+\sum_{j \neq k}^{M} \operatorname{Cov}\left(\widehat{\theta}_{j}, \widehat{\theta}_{k}\right) \phi_{j}(x) \phi_{k}(x)
\end{aligned}
$$

After integration, $\int_{0}^{1} \phi_{j}(x) \phi_{k}(x) d x=0$ so the integrated variance is

$$
\int_{0}^{1} \operatorname{Var}\left(\widehat{m}_{M}(x)\right) d x=\frac{\sigma^{2}}{n} \sum_{j=1}^{M} \int_{0}^{1} \phi_{j}^{2}(x) d x=\frac{\sigma^{2} M}{n}=O\left(\frac{M}{n}\right)
$$

Recall that the MISE is just the sum of integrated bias and integrated variance, we obtain

$$
\operatorname{MISE}\left(\widehat{m}_{M}\right)=\int_{0}^{1} \operatorname{bias}^{2}\left(\widehat{m}_{M}(x)\right) d x+\int_{0}^{1} \operatorname{Var}\left(\widehat{m}_{M}(x)\right) d x=O\left(M^{-4}\right)+O\left(\frac{M}{n}\right)
$$

Thus, the optimal choice is

$$
M^{*} \asymp n^{1 / 5}
$$

### 3.6.2 Basis approach as a linear smoother

The basis estimator is another linear smoother. To see this, we use the follow expansion:

$$
\begin{aligned}
\widehat{m}_{M}(x) & =\sum_{j=1}^{M} \widehat{\theta}_{j} \phi_{j}(x) \\
& =\sum_{j=1}^{M} \frac{1}{n} \sum_{i=1}^{n} Y_{i} \phi_{j}\left(X_{i}\right) \phi_{j}(x) \\
& =\sum_{i=1}^{n}\left(\sum_{j=1}^{M} \frac{1}{n} \phi_{j}\left(X_{i}\right) \phi_{j}(x)\right) Y_{i} \\
& =\sum_{i=1}^{n} \ell_{i}(x) Y_{i},
\end{aligned}
$$

where $\ell_{i}(x)=\sum_{j=1}^{M} \frac{1}{n} \phi_{j}\left(X_{i}\right) \phi_{j}(x)$.
Recall that from the linear smoother theory, we can estimate $\sigma^{2}$ using the residuals and the degree of freedom:

$$
\widehat{\sigma}^{2}=\frac{1}{n-2 \nu+\widetilde{\nu}} \sum_{i=1}^{n} e_{i}^{2}
$$

where $e_{i}=\widehat{Y}_{i}-Y_{i}=\widehat{m}_{M}\left(X_{i}\right)-Y_{i}$ and $\nu, \widetilde{\nu}$ are the degree of freedoms (see the previous lecture note).
With this variance estimator and the fact that $\operatorname{Var}\left(\widehat{m}_{M}(x)\right)=\frac{\sigma^{2}}{n} \sum_{j=1}^{M} \phi_{j}^{2}(x)$ and the asymptotic normality, we can construct a confidence interval (band) of $m$ using

$$
\widehat{m}_{M}(x) \pm z_{1-\alpha / 2} \frac{\widehat{\sigma}^{2}}{n} \sum_{j=1}^{M} \phi_{j}^{2}(x)
$$

Note that this confidence interval is valid for $\mathbb{E}\left(\widehat{m}_{M}(x)\right)=\sum_{j=1}^{M} \theta_{j} \phi_{j}(x)$, not the actual $m(x)$. The difference between them is the bias of our estimator.

### 3.7 Regression Tree

In this section, we assume that the covariate may have multiple dimensions, i.e., $x=\left(x_{1}, \cdots, x_{d}\right)$. And our data are $\left(X_{1}, Y_{1}\right), \cdots,\left(X_{n}, Y_{n}\right) \sim P$ for some CDF $P$. Again, we are interested in the regression function $m(x)=\mathbb{E}\left(Y_{1} \mid X_{1}=x\right)$.

Regression tree constructs an estimator of the form:

$$
m(x)=\sum_{\ell=1}^{M} c_{\ell} I\left(x \in R_{\ell}\right)
$$

where $R_{\ell}$ is some rectangle partition of the space of covariates.
Here is an example of a regression tree and its splits. In this example, there are two covariates (namely, $d=2$ ) and we have 3 regions $R_{1}, R_{2}, R_{3}$ :

$$
R_{1}=\left\{\left(x_{1}, x_{2}\right): x_{1}<10, x_{2}<5\right\}, \quad R_{2}=\left\{\left(x_{1}, x_{2}\right): x_{1}<10, x_{2} \geq 5\right\}, R_{3}=\left\{\left(x_{1}, x_{2}\right): x_{1} \geq 10\right\}
$$




A regression tree estimator will predict the same value of the response $Y$ within the same area of the covariate. Namely, $m(x)$ will be the same when $x$ is within the same area.

To use a regression tree, there are $2 M$ quantities to be determined: the regions $R_{1}, \cdots, R_{M}$ and the predicted values $c_{1}, \cdots, c_{M}$. When $R_{1}, \cdots, R_{M}$ are given, $c_{1}, \cdots, c_{M}$ can be simply estimated by the average within each region, i.e.,

$$
\widehat{c}_{\ell}=\frac{\sum_{i=1}^{n} Y_{i} I\left(X_{i} \in R_{\ell}\right)}{\sum_{i=1}^{n} I\left(X_{i} \in R_{\ell}\right)} .
$$

Thus, the difficult part is the determination of $R_{1}, \cdots, R_{M}$.
Unfortunately, there is no simple closed form solution to these regions. We only have a procedure for computing it. Here is what we will do in practice. Let $X_{i j}$ be the $j$-th coordinate of the $i$-th observation $\left(X_{i}\right)$.

1. For a given $j$, we define

$$
R_{a}(j, s)=\left\{x_{j}: x_{j}<s\right\}, \quad R_{b}(j, s)=\left\{x_{j}: x_{j} \geq s\right\} .
$$

2. Find $c_{a}$ and $c_{b}$ that minimizes

$$
\sum_{X_{i} \in R_{a}}\left(Y_{i}-c_{a}\right)^{2}, \quad \sum_{X_{i} \in R_{b}}\left(Y_{i}-c_{b}\right)^{2}
$$

respectively.
3. Compute the score

$$
S(j, s)=\sum_{X_{i} \in R_{a}}\left(Y_{i}-c_{a}\right)^{2}+\sum_{X_{i} \in R_{b}}\left(Y_{i}-c_{b}\right)^{2}
$$

4. Change $s$ and repeat the same calculation until we find the minimizer of $S(j, s)$, denoted the minimal score as $S^{*}(j)$.
5. Compute the score $S^{*}(j)$ for $j=1, \cdots, d$.
6. Pick the dimension (coordinate) and the corresponding split point $s$ that has the minimal score $S^{*}(j)$. Partition the space into two parts according to this split.
7. Repeat the above procedure for each partition until certain stopping criterion is satisfied.

Using the above procedure, we will eventually end up with a collection of rectangle partitions $\widehat{R}_{1}, \cdots, \widehat{R}_{M}$. Then the final estimator is

$$
\widehat{m}(x)=\sum_{\ell=1}^{M} \widehat{c}_{\ell} I\left(x \in \widehat{R}_{\ell}\right)
$$

For the stopping criterion, sometimes people will pick the number of $M$ so as long as we obtain $M$ regions, the splitting procedure will stop. However, such a choice $M$ is rather arbitrary. A popular alternative is to top the criterion based on minimizing some score that balances the fitting quality and the complexity of the tree. For instance, we may stop the criterion if the following score is no longer decreasing:

$$
C_{\lambda, n}(M)=\frac{1}{n} \sum_{i=1}^{n}\left(Y_{i}-\widehat{m}\left(X_{i}\right)\right)^{2}+\lambda M
$$

where $\lambda>0$ is a tuning parameter that determines the 'penalty' for having a complex tree. In the next lecture, we will talk more about this penalty type tuning parameter.

Cross-validation. The tuning parameter How to choose the tuning parameter $\lambda$ ? There is a simple approach called the cross-validation ${ }^{4}$ that can compute a good choice of this quantity. Not only $\lambda$, other tuning parameters such as the number of basis $M$, the smoothing bandwidth $h$, the bin size $b$, can be chosen using the cross-validation.

## Remark.

- Interpreation. Regression tree has a powerful feature that it is easy to interpret. Even without much training, a practitioner can use the output from a regression tree very easily. A limitation of the regression tree is that it partitions the space of covariates into rectangle regions, which may be unrealistic for the actual regression model.

[^3]- MARS (multivariate adaptive regression splines). The regression tree has another limitation that it predicts the same value within the same region. This creates a jump on the boundary of two consecutive regions. There is a modified regression tree called MARS (multivariate adaptive regression splines) that allows a continuous (and possibly smooth) changes over two regions. See https://en.wikipedia.org/wiki/Multivariate_adaptive_regression_splines.


### 3.8 Penalized regression

In the regression tree, we talk about the case that we want to select the number of leaves $M$ based on the following criterion:

$$
\begin{equation*}
C_{\lambda, n}(M)=\underbrace{\frac{1}{n} \sum_{i=1}^{n}\left(Y_{i}-\widehat{m}\left(X_{i}\right)\right)^{2}}_{\text {fitting to the data }}+\underbrace{\lambda M}_{\text {penalty on the complexity }} \tag{3.11}
\end{equation*}
$$

It turns out that this type of criterion is very general in regression analysis because we want to avoid the problem of overfitting.

The overfitting means that you fit a too complex model to the data so that although the fitted curve is close to most of the observations, the actual prediction is very bad. For instance, the following picture shows the fitted result using a smoothing/cubic spline (here the quantity spar is related to $\lambda$ ):


This data is generated from a sine function plus a small noise. When $\lambda$ is too small (orange curve), we fit a very complicated model to the data, which does not capture the right structure. On the other hand, when $\lambda$ is too large (green curve), we fit a too simple model (a straight line), which is also bad in predicting the actual outcome. When $\lambda$ is too small, it is called overfitting (orange curve) whereas when $\lambda$ is too large, it is called underfitting (green curve). In fact, overfitting is similar to undersmoothing and underfitting is similar to oversmoothing. It regression analysis, people prefer to use overfitting and underfitting to describe the outcome and in density estimation, people prefer to use undersmoothing and oversmoothing.

Finding a regression estimator using a criterion with a fitting to the data plus a penalty on the complexity is called a penalized regression. In the case of regression tree, let

$$
\mathcal{M}_{\text {Tree }}=\{\text { all possible regression trees }\}
$$

be the collection of all possible regression trees. We can rewrite equation (3.11) as

$$
\widehat{m}_{\text {Tree }}=\underset{m \in \mathcal{M}_{\text {Tree }}}{\operatorname{argmin}} \frac{1}{n} \sum_{i=1}^{n}\left(Y_{i}-m\left(X_{i}\right)\right)^{2}+\mathcal{P}_{\lambda}(m),
$$

where $\mathcal{P}_{\lambda}(m)=\lambda \times$ number of regions in $m$. Thus, with the penalty on the number of regions, the regression tree is a penalized regression approach.

For any penalized regression approach, there is an abstract expression for them:

$$
\begin{equation*}
\widehat{m}=\underset{m \in \mathcal{M}}{\operatorname{argmin}} \frac{1}{n} \sum_{i=1}^{n}\left(Y_{i}-m\left(X_{i}\right)\right)^{2}+\mathcal{P}_{\lambda}(m) \tag{3.12}
\end{equation*}
$$

where $\mathcal{M}$ is a collection of regression estimators and $\mathcal{P}_{\lambda}(m)$ is the amount of penalty imposed for a regression estimator $m \in \mathcal{M}$ and $\lambda$ is a tuning parameter that determines the amount of penalty. The penalized regression always have a fitting part (e.g., $\left.\frac{1}{n} \sum_{i=1}^{n}\left(Y_{i}-m\left(X_{i}\right)\right)^{2}\right)$ and a penalized part (also called regularized part) $\mathcal{P}_{\lambda}(m)$. The fitting part makes sure the model fits the data well while the penalized part guarantees that the model is not too complex. Thus, the penalized regression often leads to a simple model with a good fitting to the data.

### 3.9 Spline

Smoothing spline is a famous example in penalized regression methods. Here we consider the case of univariate regression (i.e., the covariate $X$ is univariate or equivalently, $d=1$ ) and focus on the region where the covariates belongs to $[0,1]$. Namely, our data is $\left(X_{1}, Y_{1}\right), \cdots,\left(X_{n}, Y_{n}\right)$ with $X_{i} \in[0,1] \subset \mathbb{R}$ for each $i$.

Let $\mathcal{M}_{2}$ denotes the collection of all univariate functions with second derivative on $[0,1]$. The cubic (smoothing) spline finds an estimator

$$
\begin{equation*}
\widehat{m}=\underset{m \in \mathcal{M}_{2}}{\operatorname{argmin}} \frac{1}{n} \sum_{i=1}^{n}\left(Y_{i}-m\left(X_{i}\right)\right)^{2}+\lambda \int_{0}^{1}\left|m^{\prime \prime}(x)\right|^{2} d x \tag{3.13}
\end{equation*}
$$

In the cubic spline the penalty function is $\lambda \int_{0}^{1}\left|m^{\prime \prime}(x)\right|^{2} d x$, which imposes restriction on the smoothness the curve $m(x)$ cannot change too drastically otherwise the second derivatives will be large. Thus, the cubic spline leads to a smooth curve but fits to the data well.

Why the estimator $\widehat{m}$ is called a cubic spline? This is because it turns out that $\widehat{m}$ is a piecewise polynomial function (spline) with degree of 3. Namely, there exists knots $\tau_{1}<\cdots<\tau_{K}$ such that for $x \in\left(\tau_{k}, \tau_{k+1}\right)$,

$$
\widehat{m}(x)=\gamma_{0, k}+\gamma_{1, k} x+\gamma_{2, k} x^{2}+\gamma_{3, k} x^{3}
$$

for some $\gamma_{0, k}, \cdots, \gamma_{3, k}$ with restriction that $\widehat{m}(x)$ has continuous second derivatives at each knot. In the case of cubic spline, it turns out that the knots are just data points.

The representation of a cubic spline is often done using some basis function. Here we will introduce a simple basis called the truncated power basis. Let $X_{(1)}<X_{(2)}<\cdots<X_{(n)}$ be the ordered statistics of $X_{1}, \cdots, X_{n}$. In the cubic spline, the knots are

$$
\tau_{1}=X_{(1)}, \tau_{2}=X_{(2)}, \cdots, \tau_{n}=X_{(n)}
$$

The truncated power basis uses a collection of functions

$$
h_{1}(x)=1, h_{2}(x)=x, h_{3}(x)=x^{2}, h_{4}(x)=x^{3},
$$

and

$$
h_{j}(x)=\left(x-\tau_{j-4}\right)_{+}^{3}, \quad j=5,6, \cdots, n+4
$$

where $(x)_{+}=\max \{x, 0\}$. Then the estimator $\widehat{m}$ can be written as

$$
\widehat{m}(x)=\sum_{j=1}^{n+4} \widehat{\beta}_{j} h_{j}(x),
$$

for some properly chosen $\widehat{\beta}_{j}$.
How do we compute $\widehat{\beta}_{1}, \cdots, \widehat{\beta}_{n+4}$ ? They should be chosen using equation (3.13). Here how we will compute it. Define an $n \times(n+4)$ matrix $\mathbb{H}$ such that

$$
\mathbb{H}_{i j}=h_{j}\left(X_{i}\right)
$$

and an $(n+4) \times(n+4)$ matrix $\Omega$ with

$$
\Omega_{i j}=\int_{0}^{1} h_{i}^{\prime \prime}(x) h_{j}^{\prime \prime}(x) d x
$$

In this case, we define $m(x)=\sum_{j=1}^{n+4} \beta_{j} h_{j}(x)$ so the criterion in the right-hand side of (3.13) becomes

$$
\begin{aligned}
\frac{1}{n} \sum_{i=1}^{n}\left(Y_{i}-m\left(X_{i}\right)\right)^{2} & +\lambda \int_{0}^{1}\left|m^{\prime \prime}(x)\right|^{2} d x \\
& =\frac{1}{n} \sum_{i=1}^{n}\left(Y_{i}-\sum_{j=1}^{n+4} \beta_{j} h_{j}\left(X_{i}\right)\right)^{2}+\lambda \int_{0}^{1}\left(\sum_{j=1}^{n+4} \beta_{j} h_{j}(x)\right)\left(\sum_{\ell=1}^{n+4} \beta_{\ell} h_{\ell}(x)\right) d x \\
& =\|\mathbb{Y}-\mathbb{H} \beta\|^{2}+\lambda \beta^{T} \Omega \beta \\
& =R_{n}(\beta)
\end{aligned}
$$

where $\mathbb{Y}=\left(Y_{1}, \cdots, Y_{n}\right)$ and $\beta=\left(\beta_{1}, \cdots, \beta_{n+4}\right)$. Thus,

$$
\widehat{\beta}=\operatorname{argmin}_{\beta} R_{n}(\beta)=\left(\mathbb{H}^{T} \mathbb{H}+\lambda \Omega\right)^{-1} \mathbb{H}^{T} \mathbb{Y}
$$

Given a point $x$, let $H(x)=\left(h_{1}(x), h_{2}(x), \cdots, h_{n+4}(x)\right)$ be an $(n+4)$-dimensional vector. Then the predicted value $\widehat{m}(x)$ has a simple form:

$$
\widehat{m}(x)=H^{T}(x) \widehat{\beta}=H^{T}(x)\left(\mathbb{H}^{T} \mathbb{H}+\lambda \Omega\right)^{-1} \mathbb{H}^{T} \mathbb{Y}=\sum_{i=1}^{n} \ell_{i}(x) Y_{i}
$$

where

$$
\ell_{i}(x)=H^{T}(x)\left(\mathbb{H}^{T} \mathbb{H}+\lambda \Omega\right)^{-1} \mathbb{H}^{T} e_{i},
$$

with $e_{i}=(0,0, \cdots, 0, \underbrace{1}_{i \text {-th coordinate }}, 0, \cdots, 0)$ is the unit vector in the $i$-th coordinate. Therefore, again the cubic spline is a linear smoother.

Note that when the sample size $n$ is large, the spline estimator behaves like a kernel regression in the sense that

$$
\ell_{i}(x) \approx \frac{1}{p\left(X_{i}\right) h\left(X_{i}\right)} K\left(\frac{X_{i}-x}{h\left(X_{i}\right)}\right)
$$

and

$$
h(x)=\left(\frac{\lambda}{n p(x)}\right)^{1 / 4}, \quad K(x)=\frac{1}{2} \exp \left(-\frac{|x|}{\sqrt{2}}\right) \sin \left(\frac{|x|}{\sqrt{2}}+\frac{\pi}{4}\right)
$$

This is formally stated in the following paper:

Silverman, B. W. (1984). Spline smoothing: the equivalent variable kernel method. The Annals of Statistics, 12(3), 898-916.

## Remark.

- Regression spline. In the case where we use the spline basis to do regression but without a penalty and use fewer number of knots (and we allow the knots to be at non data points), the resulting estimator is called a regression spline. Namely, a regression spline is an estimator of the form $\widehat{m}(x)=$ $\sum_{j=1}^{M} \widehat{\beta}_{j} h_{j}(x)$, where $\widehat{\beta}_{1}, \cdots, \widehat{\beta}_{M}$ are determined by minimizing

$$
\frac{1}{n} \sum_{i=1}^{n}\left(Y_{i}-\sum_{j=1}^{M} \beta_{j} h_{j}\left(X_{i}\right)\right)^{2}
$$

Using our notations, the regression spline can be written as

$$
\widehat{m}(x)=H^{T}(x) \widehat{\beta}=H^{T}(x)\left(\mathbb{H}^{T} \mathbb{H}\right)^{-1} \mathbb{H}^{T} \mathbb{Y}
$$

- B-spline basis. There are other basis that can be used in constructing a spline estimator. One of the most famous basis is the B-spline basis. This basis is defined through a recursive way so we will not go to the details here. If you are interested in, you can check https://cran.r-project.org/ web/packages/crs/vignettes/spline_primer.pdf. The advantage of using a B-spline basis is the computation.
- M-th order spline. There are higher order spline. If we modify the optimization criterion to

$$
\frac{1}{n} \sum_{i=1}^{n}\left(Y_{i}-m\left(X_{i}\right)\right)^{2}+\lambda \int_{0}^{1}\left|m^{(\beta)}(x)\right|^{2} d x
$$

where $m^{(\beta)}$ denotes the $\beta$-th derivative, then the estimator is called a $(\beta+1)$-th order spline. As you may expect, we can construct a truncated power basis using polynomials up to the order of $\beta+1$. Namely, we will use $1, x, x^{2}, \cdots, x^{\beta+1}$ and knots to construct the basis.


[^0]:    ${ }^{1}$ https://en.wikipedia.org/wiki/Kernel_regression

[^1]:    ${ }^{2}$ https://en.wikipedia.org/wiki/Cross-validation_(statistics)

[^2]:    ${ }^{3}$ https://en.wikipedia.org/wiki/Overfitting

[^3]:    ${ }^{4}$ https://en.wikipedia.org/wiki/Cross-validation_(statistics)

