# IEMS 304 Lecture 2: Simple Linear Regression

Yiping Lu
yiping.lu@northwestern.edu

*Industrial Engineering & Management Sciences*
*Northwestern University*

# Simple Linear Regression

# Linear Regression

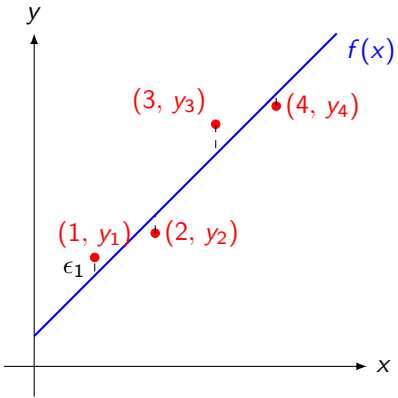


$$Y_i = \beta_0 + \beta_1 X_i + \varepsilon_i,$$

- $X$ has an arbitrary distribution, possibly deterministic.
- If $X = x$, then $Y = \beta_0 + \beta_1 x + \varepsilon$, with $\beta_0, \beta_1$ being the *coefficients*, and $\varepsilon$ being the *noise* variable.
- $\mathbb{E}[\varepsilon | X = x] = 0,\ \mathrm{Var}(\varepsilon | X = x) = \sigma^2$.

# Least Squares Estimators

One option to estimate the unknown quantities is to find the optimal fit, to be precise here, minimize the mean squared error (MSE):

$$(\beta_0, \beta_1) = \arg \min_{(b_0, b_1)} \mathbb{E}[(Y - (b_0 + b_1 X)^2)].$$

- How to access $\mathbb{E}$?
  - The data we may consider are $\{(X_1, Y_1), \ldots, (X_n, Y_n)\}$.

## How to Estimate $\pi$ ?

- Draw a square of side length 2 (from $-1$ to $+1$) and inscribe a circle of radius 1.
- Randomly sample the points within the square.
- Count how many points fall inside the circle.
- The expectation of fraction of points in the circle is $\frac{\text{the circle's area}}{\text{total points' area}} \approx \frac{\pi r^2}{(2r)^2} = \frac{\pi}{4}$.
- Hence $\pi \approx 4 \times \frac{\text{points in circle}}{\text{total points}}$.

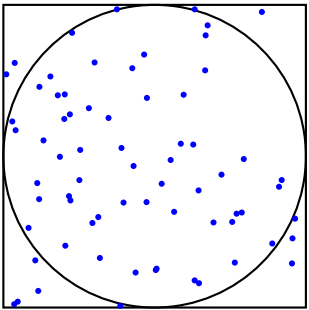

# Find $\beta_0, \beta_1$

We minimize in-sample, empirical MSE:

$$(\hat{\beta}_0, \hat{\beta}_1) = \arg \min_{(b_0, b_1)} \underbrace{\frac{1}{n} \sum_{i=1}^{n} (Y_i - (b_0 + b_1 X_i))^2}_{\widehat{\mathrm{MSE}}(b_0, b_1)}.$$

---

**Next.** $\hat{\beta}_0, \hat{\beta}_1$ has closed form solution!

**How ?**

# How to find the Minimizer of a Function

How to find the Minimizer of a function $x^* = \arg\min_x f(x)$?

Solve the equation $\nabla f(x^*) = 0$

# Find $\beta_0, \beta_1$

$$\hat{\beta}_1 = \frac{c_{XY}}{s_X^2},$$

where $c_{XY}, s_X^2$ are the sample covariance between $X, Y$ and the sample variance of $X$ respectively. As a reminder,

$$c_{XY} = \frac{1}{n}\sum_{i=1}^{n}(X_i - \overline{x})(Y_i - \overline{y}), s_X^2 = \frac{1}{n}\sum_{i=1}^{n}(X_i - \overline{x})^2.$$

---

$$0 = \overline{xy} - (\overline{y} - \hat{\beta}_1\overline{x})\overline{x} - \hat{\beta}_1\overline{x^2}$$

$$0 = c_{XY} - \hat{\beta}_1 s_X^2$$

# How accurate is the Model?– Bias

$$\hat{\beta}_1 = \beta_1 + \frac{1}{n s_X^2} \sum_{i=1}^{n} (X_i - \bar{x}) \varepsilon_i.$$

---

**Statement:** $\hat{\beta}_1$ is unbiased, i.e. $\mathbb{E}[\hat{\beta}_1] = \beta_1$.

# Model Fitting

- Find $(\hat{\beta}_0, \hat{\beta}_1)$ that minimize the least square

$$Q = \sum_{i=1}^{n} (y_i - \underbrace{(\hat{\beta}_0 + \hat{\beta}_1 x_i)}_{\hat{y}_i})^2.$$

  - Denote $\hat{y}_i = \hat{\beta}_0 + \hat{\beta}_1 x_i$ as the **fitted value**;
  - Denote $e_i = y_i - \hat{y}_i$ as the **residual**.

  Therefore, minimizing the least square can be understood as fitting $y_i$'s to minimize residuals as good as possible.

# How accurate is the Model?– Variance

$$\mathrm{Var}(\hat{\beta}_1) = \mathrm{Var}\left(\beta_1 + \frac{1}{ns_X^2}\sum_{i=1}^{n}(X_i - \overline{x})\varepsilon_i\right) = \frac{\sigma^2}{ns_X^2}.$$

# Unconditioning on $X$

- **Bias** apply the law of total expectation:

$$\mathbb{E}[\hat{\beta}_1] = \mathbb{E}\Big[\mathbb{E}[\hat{\beta}_1 \mid X_1, \ldots, X_n]\Big] = \mathbb{E}[\beta_1] = \beta_1.$$

- **Variance** apply the law of total variance:

$$\begin{aligned}\mathrm{Var}(\hat{\beta}_1) &= \mathbb{E}\Big[\mathrm{Var}(\hat{\beta}_1 \mid X_1, \ldots, X_n)\Big] + \mathrm{Var}\Big(\mathbb{E}[\hat{\beta}_1 \mid X_1, \ldots, X_n]\Big) \\ &= \mathbb{E}\left[\frac{\sigma^2}{n s_X^2}\right] + \mathrm{Var}(\beta_1) = \frac{\sigma^2}{n}\mathbb{E}\left[\frac{1}{s_X^2}\right].\end{aligned}$$

# Go Beyond Point Estimation

**Fact.** $\mathbb{E}[\hat{f}(x)] = \beta_0 + \beta_1 x$. and $\mathrm{Var}(\hat{f}(x)) = \frac{\sigma^2}{n}\left(1 + \frac{(x-\bar{x})^2}{s_X^2}\right)$.

What is the the standard error of an estimator ? $\mathrm{se}(\hat{\beta}_1) = \frac{\sigma}{\sqrt{n s_X^2}}$.

# Exercise

- What happens when the noise variance, $\sigma^2$, increases?
- What happens when the number of samples, $n$, increases?
- What influences the variance of our predictions?
- What happens when we predict at $x$ that is very close to $\overline{x}$? How about very far?

# How to Estimate $\sigma$?

Using the simple linear regression model,

$$\mathbb{E}[(Y - (\beta_0 + \beta_1 X))^2] = \sigma^2. \quad \textit{(convince yourself why.)}$$

Then, a natural estimator for $\sigma^2$ would be

$$\hat{\sigma}^2 = \frac{1}{n}\sum_{i=1}^{n}(Y_i - \hat{f}(X_i))^2.$$

Notice that this is a biased estimator. Moreover $s^2 = \frac{n}{n-2}\hat{\sigma}^2$ is an unbiased estimator of $\sigma^2$. (Later)

# Residual and Error

$$\text{(residual)} \quad e_i = Y_i - (\hat{\beta}_0 + \hat{\beta}_1 X_i)$$

$$\text{(noise)} \quad \varepsilon_i = Y_i - (\beta_0 + \beta_1 X_i)$$

# Remark

- The sum of noise variables cannot equal zero all the time, because $\mathrm{Var}(\sum_{i=1}^{n} \varepsilon_i) = n\sigma^2$.
- The sum of residuals is *always* zero, i.e. $\sum_{i=1}^{n} e_i = 0$.
- The sample correlation between the residuals and $X_i$'s is also 0, i.e. $\sum_{i=1}^{n}(X_i - \overline{x})e_i = 0$.

# Assessing the Fit

# Assessing the Fit

- As in simple regression, we calculate
  - fitted values: $\hat{y}_i = \hat{\beta}_0 + \hat{\beta}_1 x_i$;
  - residuals: $e_i = y_i - \hat{y}_i$;
  - error sum of squares: $\mathrm{SSE} = \sum_{i=1}^n e_i^2$;
  - total sum of squares: $\mathrm{SST} = \sum_{i=1}^n (y_i - \bar{y})^2$;
  - regression sum of squares: $\mathrm{SSR} = \sum_{i=1}^n (\hat{y}_i - \bar{y})^2$.

$\bar{y} = \arg\min_c \sum_{i=1}^n (c - y_i)^2$ is the best constant fit of $\{y_i\}_{i=1}^n$!

- We can decompose $\mathrm{SST}$ as

$$\underbrace{\sum_{i=1}^n (y_i - \bar{y})^2}_{\mathrm{SST}} = \underbrace{\sum_{i=1}^n (\hat{y}_i - \bar{y})^2}_{\mathrm{SSR}} + \underbrace{\sum_{i=1}^n (y_i - \hat{y}_i)^2}_{\mathrm{SSE}}$$

# $R^2$ Statistics and Correlation

**$R^2$ (Coefficient of Determination):**

$$R^2 = \frac{\text{SSR}}{\text{SST}}, \quad \text{where} \quad \text{SSR} = \sum(\hat{y}_i - \bar{y})^2, \quad \text{SST} = \sum(y_i - \bar{y})^2.$$

**Theorem**

*Recall* *Pearson correlation coefficient:* $r = \frac{\sum(x_i - \bar{x})(y_i - \bar{y})}{\sqrt{\sum(x_i - \bar{x})^2 \sum(y_i - \bar{y})^2}}$, *then we have*

$$R^2 = r^2$$

# Prove $R^2 = r^2$

Since $\hat{\beta}_1 = \frac{\sum(x_i-\bar{x})(y_i-\bar{y})}{\sum(x_i-\bar{x})^2} = r\,\frac{s_y}{s_x}$, we have $\text{SSR} = \frac{(\sum(x_i-\bar{x})(y_i-\bar{y}))^2}{\sum(x_i-\bar{x})^2}$. Thus,

$$R^2 = \frac{\text{SSR}}{\text{SST}} = \frac{\left(\sum(x_i-\bar{x})(y_i-\bar{y})\right)^2}{\sum(x_i-\bar{x})^2\ \sum(y_i-\bar{y})^2} = r^2.$$

# Error

**Prove**: $s^2 = \frac{n}{n-2}\hat{\sigma}^2$ is an *unbiased* estimator of $\sigma^2$

# Pipeline of Machine Learning

# Log-Likelihood

The model looks similar,

$$Y_i = \beta_0 + \beta_1 X_i + \varepsilon_i,$$

with modified assumptions:

- $X$ has an arbitrary distribution, possibly deterministic.
- If $X = x$, then $Y = \beta_0 + \beta_1 x + \varepsilon$, with $\beta_0, \beta_1$ being the coefficients, and $\varepsilon$ being the noise variable.
- (stronger) $\varepsilon \sim N(0, \sigma^2)$, and is independent of $X$.
- (stronger) $\varepsilon$ is *independent* across observations.

---

**Question.** What is $p(Y_i | X_i; b_0, b_1, s^2)$?

# Log-Likelihood

Given the data, the likelihood under this set of assumption is a function of the unknown parameters, defined as

$$L(b_0, b_1, s^2) = \prod_{i=1}^{n} p(Y_i|X_i; b_0, b_1, s^2) = \prod_{i=1}^{n} \frac{1}{\sqrt{2\pi s^2}} \exp\left\{ -\frac{1}{2s^2}(Y_i - (b_0 + b_1 X_i))^2 \right\}.$$

$$\log(ab) = \log(a) + \log(b)$$

---

$$\log L(b_0, b_1, s^2) \stackrel{\text{def}}{=} \ell(b_0, b_1, s^2) = -\frac{n}{2}\log(2\pi) - \frac{n}{2}\log s^2 - \frac{1}{2s^2}(Y_i - (b_0 + b_1 X_i))^2.$$

# Logistic regression

**Step 1. Likelihood for a Logistic Binary Outcome:**
For each observation $y_i \in \{0, 1\}$ with probability $p_i$ for $y_i = 1$, the likelihood is

$$L(p_i \mid y_i) = p_i^{y_i}(1 - p_i)^{1-y_i}.$$

where probability $p_i = \frac{1}{1+e^{-\beta^T x_i}}$ using the logistic function.

**Step 2. Log-Likelihood:**
For $n$ independent observations, the log-likelihood function is

$$\ell(\beta) = \sum_{i=1}^{n} \left[ y_i \log\left( \frac{1}{1 + e^{-\beta^T x_i}} \right) + (1 - y_i) \log\left( 1 - \frac{1}{1 + e^{-\beta^T x_i}} \right) \right].$$

**Step 3. Estimation:**
Maximizing $\ell(\beta)$ with respect to $\beta$ gives the maximum likelihood estimates, leading to the logistic regression model.

☹ No closed-form solution.

# Gradient Descent

- **Gradient Descent** is an iterative optimization method to find local minima of a function.
- The update rule is $x_{n+1} = x_n - \alpha \nabla f(x_n)$, where $\alpha$ is the learning rate.

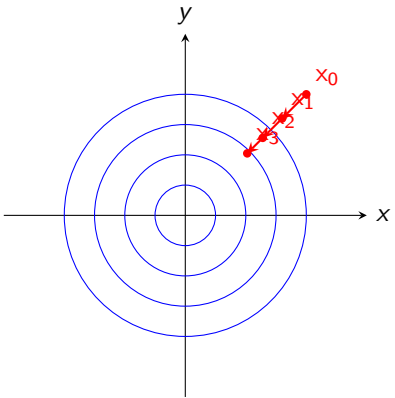

# Ill Conditioned Problems

- The function $f(x_1, x_2) = 10x_1^2 + x_2^2$ has very different curvatures along $x_1$ and $x_2$.
- Its level sets are ellipses elongated along the $x_2$-axis.
- With a fixed learning rate, gradient descent can overshoot in the steep $x_1$ direction, leading to oscillatory (zigzag) behavior.

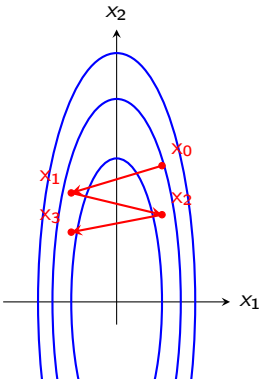

# Newton Methods

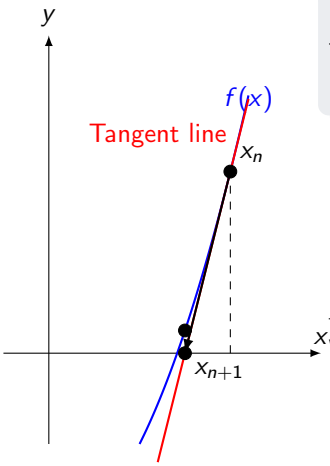


Newton's method is an iterative technique for finding a root of a nonlinear equation $F(x) = 0$ via

$$x_{n+1} = x_n - F'(x_n)^{-1} F(x_n).$$

---

What happens if one optimize $f(x_1, x_2) = 10x_1^2 + x_2^2$?

# Homework

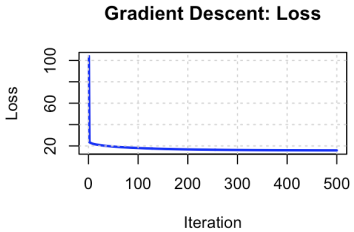
Gradient Descent: Loss
Loss
20
60
100
0
100
200
300
400
500
Iteration

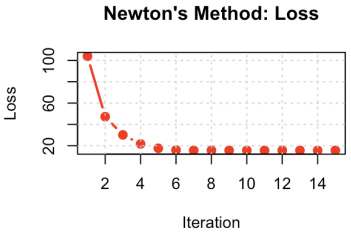
Newton's Method: Loss
Loss
20
60
100
2
4
6
8
10
12
14
Iteration

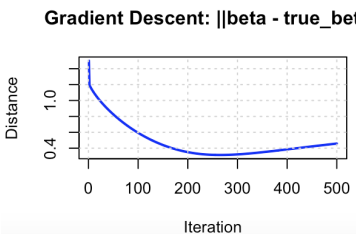
Gradient Descent: ||beta - true_beta||
Distance
0.4
1.0
0
100
200
300
400
500
Iteration

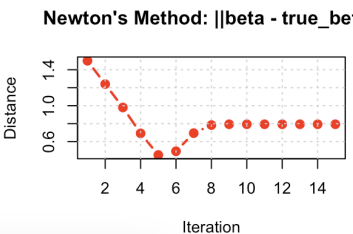
Newton's Method: ||beta - true_beta||
Distance
0.6
1.0
1.4
2
4
6
8
10
12
14
Iteration

# Pipeline of Machine Learning