COMS 4771 Lecture 15

1. Ordinary least squares (three perspectives)
2. Computation

Regression Analysis

So far, mostly looked at **classification** problems, where the label space $\mathcal{Y} = \{1, 2, \ldots, K\}$ is discrete.

Regression Analysis

So far, mostly looked at **classification** problems, where the label space $\mathcal{Y} = \{1, 2, \ldots, K\}$ is discrete.

Regression analysis considers prediction problems where label space $\mathcal{Y}$ is continuous (e.g. $[0, 1]$, $\mathbb{R}^d$, etc.)

Regression Analysis

So far, mostly looked at **classification** problems, where the label space $\mathcal{Y} = \{1, 2, \ldots, K\}$ is discrete.

Regression analysis considers prediction problems where label space $\mathcal{Y}$ is continuous (e.g. $[0, 1]$, $\mathbb{R}^d$, etc.)

Measuring quality of a predictor $f : \mathcal{X} \to \mathcal{Y}$ using prediction error, $\Pr[f(X) \neq Y]$ doesn't make much sense in the regression context (why?)

Regression Analysis

So far, mostly looked at **classification** problems, where the label space $\mathcal{Y} = \{1, 2, \ldots, K\}$ is discrete.

Regression analysis considers prediction problems where label space $\mathcal{Y}$ is continuous (e.g. $[0, 1]$, $\mathbb{R}^d$, etc.)

Measuring quality of a predictor $f : \mathcal{X} \to \mathcal{Y}$ using prediction error, $\Pr[f(X) \neq Y]$ doesn't make much sense in the regression context (why?)

Goal: find $f : \mathcal{X} \to \mathcal{Y}$, from some function class $\mathcal{F}$, minimizing error measured using a *loss function* $\ell : \mathcal{Y} \times \mathcal{Y} \to \mathbb{R}$, i.e.

$$\mathbb{E}[\ell(Y, f(X))]$$

Regression Analysis

So far, mostly looked at **classification** problems, where the label space $\mathcal{Y} = \{1, 2, \ldots, K\}$ is discrete.

Regression analysis considers prediction problems where label space $\mathcal{Y}$ is continuous (e.g. $[0, 1]$, $\mathbb{R}^d$, etc.)

Measuring quality of a predictor $f : \mathcal{X} \to \mathcal{Y}$ using prediction error, $\Pr[f(X) \neq Y]$ doesn't make much sense in the regression context (why?)

Goal: find $f : \mathcal{X} \to \mathcal{Y}$, from some function class $\mathcal{F}$, minimizing error measured using a *loss function* $\ell : \mathcal{Y} \times \mathcal{Y} \to \mathbb{R}$, i.e.

$$\mathbb{E}[\ell(Y, f(X))]$$

E.g. if $\mathcal{Y} = \mathbb{R}$, may use $\ell(y, \hat{y}) = (y - \hat{y})^2$ ("square loss") or $\ell(y, \hat{y}) = |y - \hat{y}|$ ("absolute loss")

Linear regression (introduction)

Example: Old Faithful Geyser (Yellowstone)

Time between eruptions seems to be related to duration of previous eruption.

Example: Old Faithful Geyser (Yellowstone)

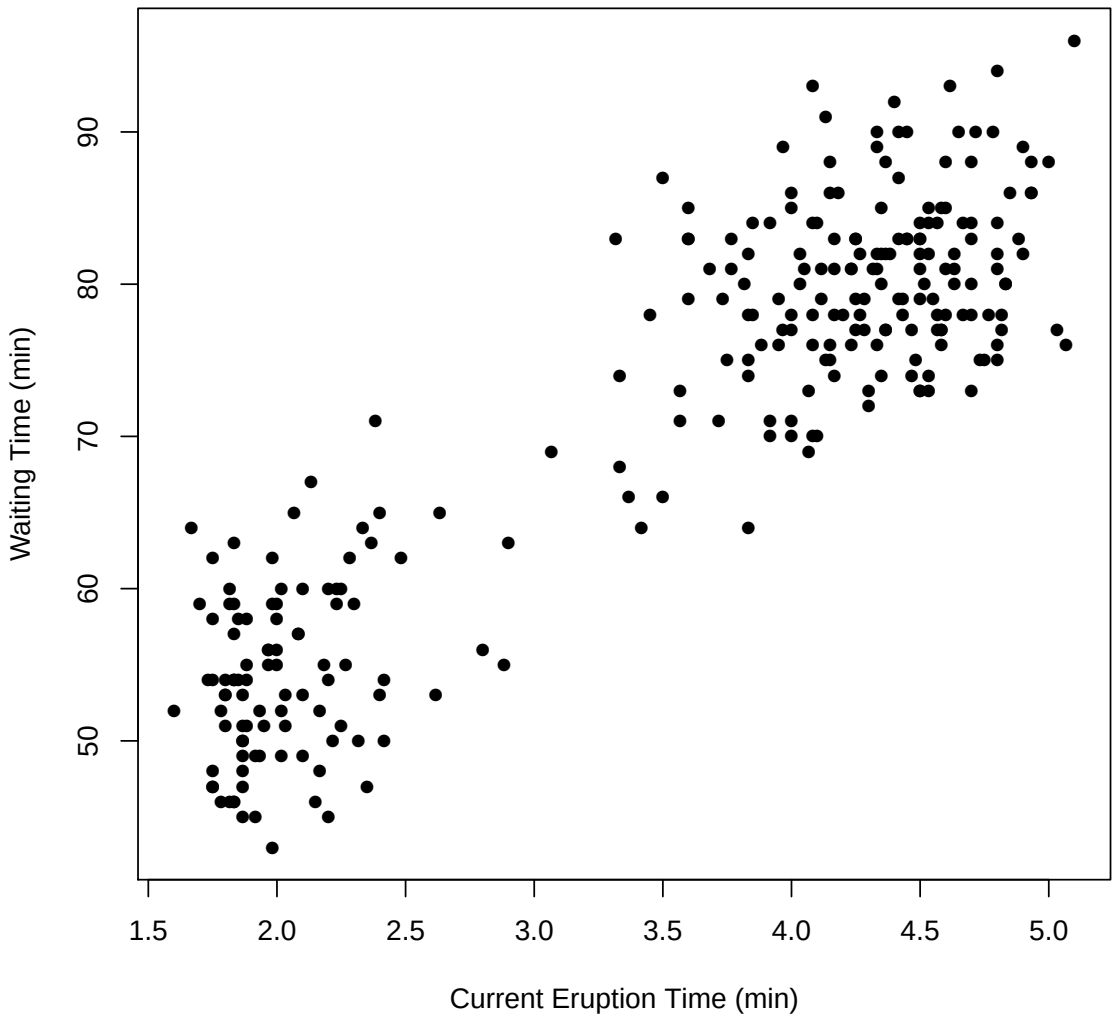

Example: Old Faithful Geyser (Yellowstone)

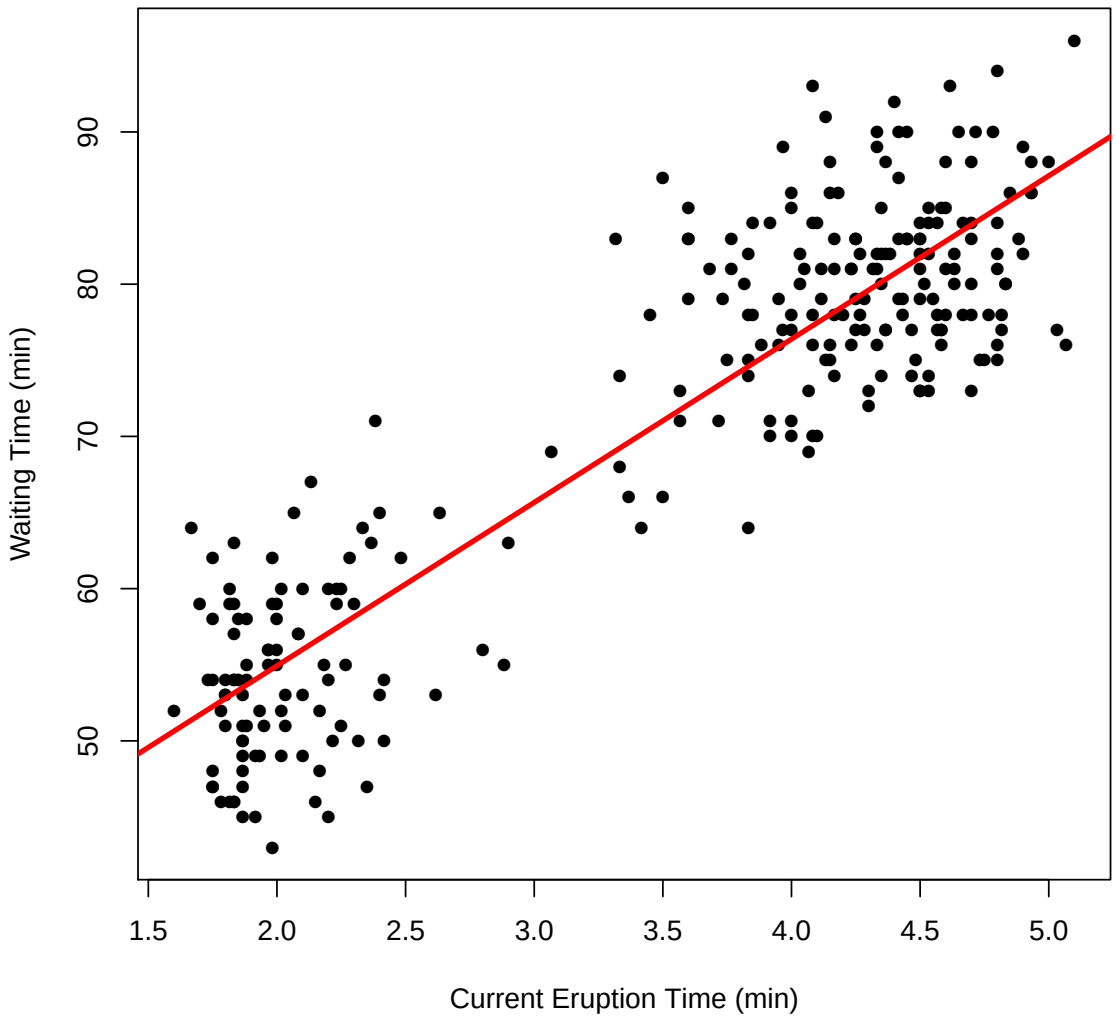

Example: Old Faithful

Linear regression model

$$\text{(wait time)} = w_0 + \text{(last duration)} \times w_1 + \text{(error)}$$

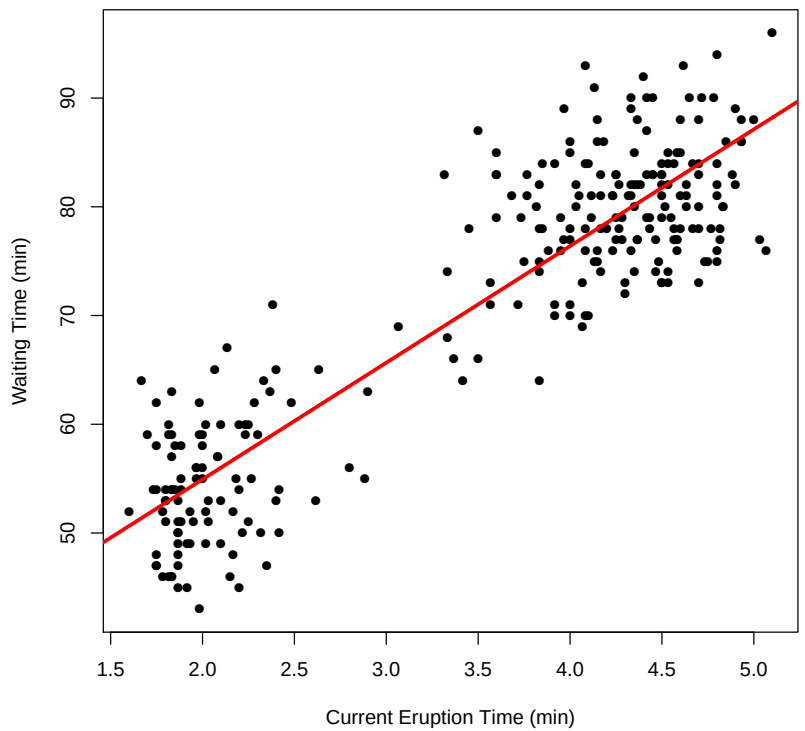

Linear regression model in $\mathbb{R}^p$

- Input variables $\boldsymbol{x} := (x_1, x_2, \ldots, x_p)$ ("covariates").
- Output variable y ("response").
- Regression coefficients $\boldsymbol{w} := (w_1, w_2, \ldots, w_p)$, intercept term w_0.

Modeling equation:

$$y = w_0 + \langle \boldsymbol{x}, \boldsymbol{w} \rangle + \varepsilon$$

where $\varepsilon := y - (w_0 + \langle \boldsymbol{x}, \boldsymbol{w} \rangle)$ is the error term.

Linear regression model in $\mathbb{R}^p$

- Input variables $\boldsymbol{x} := (x_1, x_2, \ldots, x_p)$ ("covariates").
- Output variable y ("response").
- Regression coefficients $\boldsymbol{w} := (w_1, w_2, \ldots, w_p)$, intercept term w_0.

Modeling equation:

$$y = w_0 + \langle \boldsymbol{x}, \boldsymbol{w} \rangle + \varepsilon$$

where $\varepsilon := y - (w_0 + \langle \boldsymbol{x}, \boldsymbol{w} \rangle)$ is the error term.

Least squares criterion

Given pairs of input/output values, find $(w_0, \boldsymbol{w})$ to minimize ε^2 (on average).

I.e. using square loss: $\varepsilon^2 = (y - \underbrace{(w_0 + \langle \boldsymbol{x}, \boldsymbol{w} \rangle)}_{\hat{y}})^2$.

Least squares in pictures

Red dots: data points.

$(w_0, w_1, w_2) \rightarrow$ affine hyperplane.

Vertical length is error.

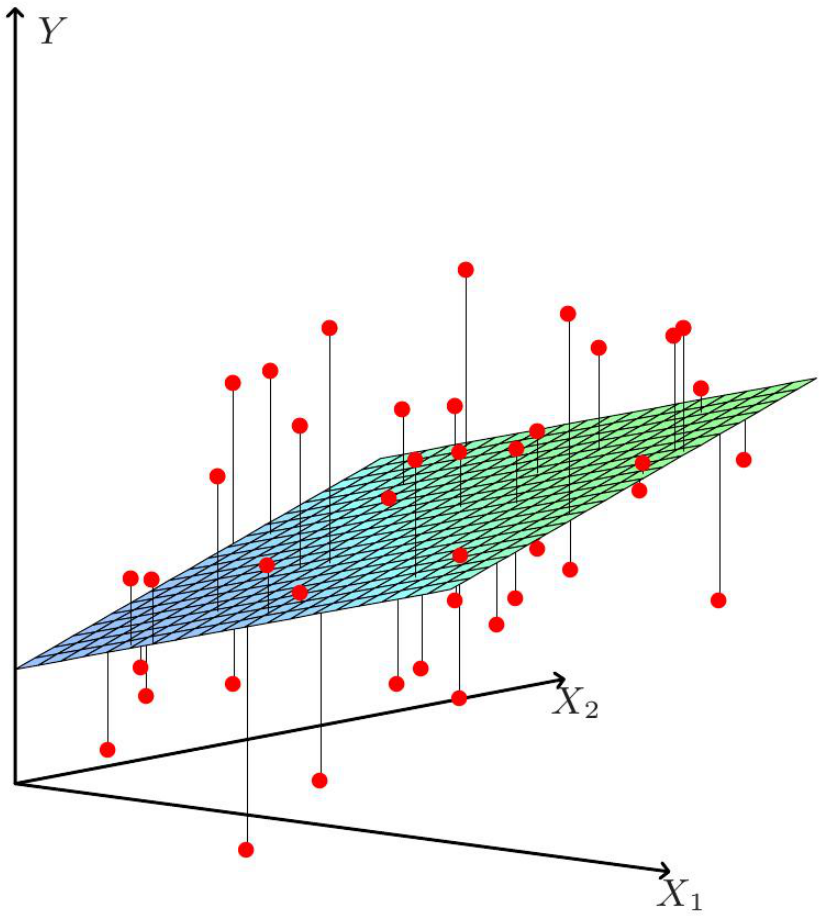

Least squares criterion

Given training data

$$\boldsymbol{X} = \begin{bmatrix} — & \boldsymbol{x}^{(1)\top} & — \\ — & \boldsymbol{x}^{(2)\top} & — \\ & \vdots & \\ — & \boldsymbol{x}^{(n)\top} & — \end{bmatrix} \in \mathbb{R}^{n \times p}, \qquad \boldsymbol{y} = \begin{bmatrix} y^{(1)} \\ y^{(2)} \\ \vdots \\ y^{(n)} \end{bmatrix},$$

find $w_0 \in \mathbb{R}$ and $\boldsymbol{w} \in \mathbb{R}^p$ to minimize

$$f_{\text{ls}}(\boldsymbol{w}) := \frac{1}{n} \sum_{i=1}^{n} \left(y^{(i)} - \left(w_0 + \langle \boldsymbol{x}^{(i)}, \boldsymbol{w} \rangle \right) \right)^2 = \frac{1}{n} \left\| \boldsymbol{y} - \begin{bmatrix} \mathbf{1} & \boldsymbol{X} \end{bmatrix} \begin{bmatrix} w_0 \\ \boldsymbol{w} \end{bmatrix} \right\|_2^2.$$

Simplification

Replace $\boldsymbol{X}$ with $\begin{bmatrix} \mathbf{1} & \boldsymbol{X} \end{bmatrix}$ and $\boldsymbol{w}$ with $(w_0, \boldsymbol{w})$, so least squares criterion is more simply written as

$$f_{\text{ls}}(\boldsymbol{w}) = \frac{1}{n} \|\boldsymbol{y} - \boldsymbol{X}\boldsymbol{w}\|_2^2.$$

Least squares via calculus

Least squares criterion is convex function of $\boldsymbol{w}$;
so suffices to find $\boldsymbol{w}$ where gradient is zero.

Least squares via calculus

Least squares criterion is convex function of $\boldsymbol{w}$;
so suffices to find $\boldsymbol{w}$ where gradient is zero.

Take gradient with respect to $\boldsymbol{w}$:

$$\nabla_{\boldsymbol{w}} f_{\text{ls}}(\boldsymbol{w}) = \nabla_{\boldsymbol{w}} \left\{ \frac{1}{n} \|\boldsymbol{y} - \boldsymbol{X}\boldsymbol{w}\|_2^2 \right\} = \frac{2}{n} \boldsymbol{X}^\top (\boldsymbol{X}\boldsymbol{w} - \boldsymbol{y}).$$

Least squares via calculus

Least squares criterion is convex function of $\boldsymbol{w}$;
so suffices to find $\boldsymbol{w}$ where gradient is zero.

Take gradient with respect to $\boldsymbol{w}$:

$$\nabla_{\boldsymbol{w}} f_{\text{ls}}(\boldsymbol{w}) = \nabla_{\boldsymbol{w}} \left\{ \frac{1}{n} \|\boldsymbol{y} - \boldsymbol{X}\boldsymbol{w}\|_2^2 \right\} = \frac{2}{n} \boldsymbol{X}^\top (\boldsymbol{X}\boldsymbol{w} - \boldsymbol{y}).$$

This is zero when

$$(\boldsymbol{X}^\top \boldsymbol{X})\boldsymbol{w} = \boldsymbol{X}^\top \boldsymbol{y},$$

a linear system of equations in $\boldsymbol{w}$ ("normal equations").

Least squares via calculus

Least squares criterion is convex function of $\boldsymbol{w}$;
so suffices to find $\boldsymbol{w}$ where gradient is zero.

Take gradient with respect to $\boldsymbol{w}$:

$$\nabla_{\boldsymbol{w}} f_{\text{ls}}(\boldsymbol{w}) = \nabla_{\boldsymbol{w}} \left\{ \frac{1}{n} \|\boldsymbol{y} - \boldsymbol{X}\boldsymbol{w}\|_2^2 \right\} = \frac{2}{n} \boldsymbol{X}^\top (\boldsymbol{X}\boldsymbol{w} - \boldsymbol{y}).$$

This is zero when

$$(\boldsymbol{X}^\top \boldsymbol{X})\boldsymbol{w} = \boldsymbol{X}^\top \boldsymbol{y},$$

a linear system of equations in $\boldsymbol{w}$ ("normal equations").

If $\boldsymbol{X}^\top \boldsymbol{X}$ is invertible, solution is

$$\hat{\boldsymbol{w}}_{\text{ols}} := (\boldsymbol{X}^\top \boldsymbol{X})^{-1} \boldsymbol{X}^\top \boldsymbol{y}$$

("ordinary least squares").

Column view of ordinary least squares

Column view of ordinary least squares

Least squares criterion

Let $\boldsymbol{x}_j \in \mathbb{R}^n$ be the j-th column of $\boldsymbol{X} \in \mathbb{R}^{n \times p}$, so

$$\boldsymbol{X} = \begin{bmatrix} \boldsymbol{x}_1 & \boldsymbol{x}_2 & \cdots & \boldsymbol{x}_p \end{bmatrix}.$$

Least squares criterion

Let $\boldsymbol{x}_j \in \mathbb{R}^n$ be the j-th column of $\boldsymbol{X} \in \mathbb{R}^{n \times p}$, so

$$\boldsymbol{X} = \begin{bmatrix} \boldsymbol{x}_1 & \boldsymbol{x}_2 & \cdots & \boldsymbol{x}_p \end{bmatrix}.$$

Find linear combination $\sum_{j=1}^{p} w_j \boldsymbol{x}_j$ of $\boldsymbol{x}_1, \boldsymbol{x}_2, \ldots, \boldsymbol{x}_p$ so as to minimize

$$f_{\text{ls}}(\boldsymbol{w}) \;=\; \frac{1}{n}\|\boldsymbol{y} - \boldsymbol{X}\boldsymbol{w}\|_2^2 \;=\; \frac{1}{n}\left\|\boldsymbol{y} - \sum_{j=1}^{p} w_j \boldsymbol{x}_j\right\|_2^2.$$

Least squares criterion

Let $\boldsymbol{x}_j \in \mathbb{R}^n$ be the j-th column of $\boldsymbol{X} \in \mathbb{R}^{n \times p}$, so

$$\boldsymbol{X} = \begin{bmatrix} \boldsymbol{x}_1 & \boldsymbol{x}_2 & \cdots & \boldsymbol{x}_p \end{bmatrix}.$$

Find linear combination $\sum_{j=1}^p w_j \boldsymbol{x}_j$ of $\boldsymbol{x}_1, \boldsymbol{x}_2, \ldots, \boldsymbol{x}_p$ so as to minimize

$$f_{\text{ls}}(\boldsymbol{w}) \;=\; \frac{1}{n}\|\boldsymbol{y} - \boldsymbol{X}\boldsymbol{w}\|_2^2 \;=\; \frac{1}{n}\left\|\boldsymbol{y} - \sum_{j=1}^p w_j \boldsymbol{x}_j\right\|_2^2.$$

Approximation of $\boldsymbol{y}$ via ordinary least squares (assuming $\boldsymbol{X}^\top \boldsymbol{X}$ invertible):

$$\hat{\boldsymbol{y}} = \boldsymbol{X}\hat{\boldsymbol{w}}_{\text{ols}} = \sum_{j=1}^p \hat{w}_{\text{ols},j} \boldsymbol{x}_j.$$

Column view of ordinary least squares

$\hat{\boldsymbol{y}} = \boldsymbol{X}\hat{\boldsymbol{w}}_{\text{ols}}$ is the orthogonal projection of $\boldsymbol{y}$ onto $\text{span}(\boldsymbol{x}_1, \boldsymbol{x}_2, \ldots, \boldsymbol{x}_p)$:

$$\hat{\boldsymbol{y}} = \boldsymbol{X}\hat{\boldsymbol{w}}_{\text{ols}} = \underbrace{\boldsymbol{X}(\boldsymbol{X}^\top \boldsymbol{X})^{-1}\boldsymbol{X}^\top}_{\boldsymbol{\Pi}} \boldsymbol{y}.$$

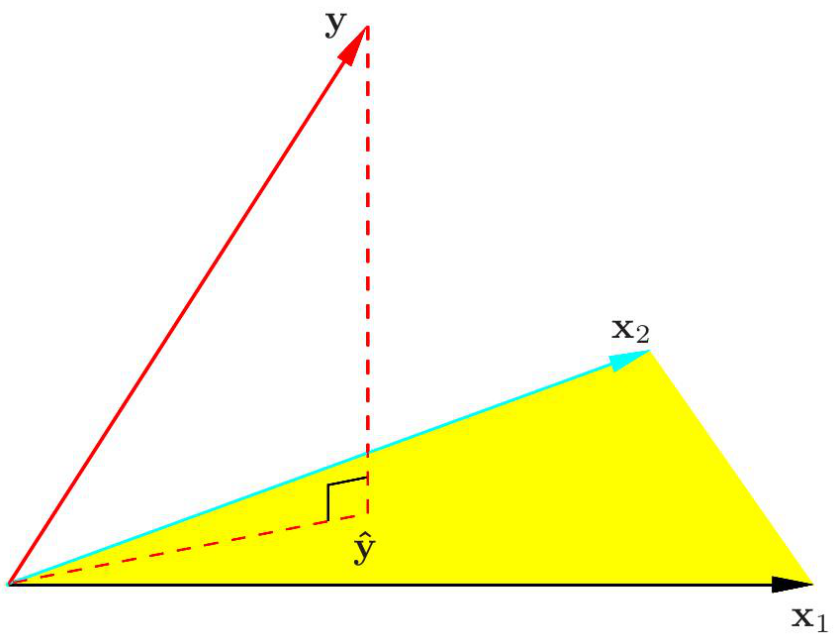

$\boldsymbol{\Pi} \in \mathbb{R}^{n \times n}$ is the orthogonal projection operator for $\text{span}(\boldsymbol{x}_1, \boldsymbol{x}_2, \ldots, \boldsymbol{x}_p)$.

Residual $\boldsymbol{r} := \boldsymbol{y} - \boldsymbol{X}\hat{\boldsymbol{w}}_{\text{ols}}$ is orthogonal to all $\boldsymbol{x}_j$.

Statistical learning perspective

Linear regression model

- Let P be a distribution over $\mathbb{R}^p \times \mathbb{R}$, and $(\boldsymbol{x}, y) \sim P$.

  Define

$$\boldsymbol{w}_\star := \operatorname*{arg\,min}_{\boldsymbol{w} \in \mathbb{R}^p} \mathbb{E}\left[\left(y - \langle \boldsymbol{x}, \boldsymbol{w} \rangle\right)^2\right]$$

Linear regression model

- Let P be a distribution over $\mathbb{R}^p \times \mathbb{R}$, and $(\boldsymbol{x}, y) \sim P$.

  Define

$$\boldsymbol{w}_\star := \operatorname*{arg\,min}_{\boldsymbol{w} \in \mathbb{R}^p} \mathbb{E}\left[\left(y - \langle \boldsymbol{x}, \boldsymbol{w} \rangle\right)^2\right]$$

- **Goal**: given i.i.d. sample S from P, find $\boldsymbol{w} \in \mathbb{R}^p$ so that execess mean squared error

$$\mathbb{E}\left[\left(y - \langle \boldsymbol{x}, \boldsymbol{w} \rangle\right)^2\right] - \mathbb{E}\left[\left(y - \langle \boldsymbol{x}, \boldsymbol{w}_\star \rangle\right)^2\right]$$

  is small (and $\to 0$ as sample size $\to \infty$).

Ordinary least squares

Ordinary least squares picks $\boldsymbol{w}$ to minimize empirical mean squared error based on i.i.d. sample S:

$$\hat{\boldsymbol{w}}_{\mathrm{ols}} := \operatorname*{arg\,min}_{\boldsymbol{w} \in \mathbb{R}^p} \sum_{(\boldsymbol{x}, y) \in S} \Big(y - \langle \boldsymbol{x}, \boldsymbol{w} \rangle \Big)^2 .$$

Ordinary least squares

Ordinary least squares picks $\boldsymbol{w}$ to minimize empirical mean squared error based on i.i.d. sample S:

$$\hat{\boldsymbol{w}}_{\mathrm{ols}} := \operatorname*{arg\,min}_{\boldsymbol{w} \in \mathbb{R}^p} \sum_{(\boldsymbol{x}, y) \in S} \Big(y - \langle \boldsymbol{x}, \boldsymbol{w} \rangle \Big)^2 .$$

- Unconstrained convex optimization problem that happens to have "closed form" solution.

Ordinary least squares

Ordinary least squares picks $\boldsymbol{w}$ to minimize empirical mean squared error based on i.i.d. sample S:

$$\hat{\boldsymbol{w}}_{\mathrm{ols}} := \operatorname*{arg\,min}_{\boldsymbol{w} \in \mathbb{R}^p} \sum_{(\boldsymbol{x}, y) \in S} \Big(y - \langle \boldsymbol{x}, \boldsymbol{w} \rangle \Big)^2 .$$

- Unconstrained convex optimization problem that happens to have "closed form" solution.
- Solution not uniquely defined unless $\sum_{(\boldsymbol{x}, y) \in S} \boldsymbol{x}\boldsymbol{x}^\top$ is invertible.

Ordinary least squares

Ordinary least squares picks $\boldsymbol{w}$ to minimize empirical mean squared error based on i.i.d. sample S:

$$\hat{\boldsymbol{w}}_{\text{ols}} := \operatorname*{arg\,min}_{\boldsymbol{w} \in \mathbb{R}^p} \sum_{(\boldsymbol{x}, y) \in S} \Big(y - \langle \boldsymbol{x}, \boldsymbol{w} \rangle \Big)^2.$$

- Unconstrained convex optimization problem that happens to have "closed form" solution.
- Solution not uniquely defined unless $\sum_{(\boldsymbol{x}, y) \in S} \boldsymbol{x}\boldsymbol{x}^\top$ is invertible.
- **Predictive performance**: (with $n = |S|$)
  - $n < p$: Could be rubbish.
  - $n \geq p$: Excess mean squared error decreases at a rate of $O\Big(\frac{p}{n}\Big)$ (under some general conditions).

Statistical estimation perspective

Maximum likelihood interpretation

Suppose the distribution P of $(\boldsymbol{x}, y)$ is such that, conditioned on $\boldsymbol{x}$,

$$y|\boldsymbol{x} \sim \mathrm{N}\big(\langle \boldsymbol{x}, \boldsymbol{w}_\star \rangle, \sigma^2\big).$$

Maximum likelihood interpretation

Suppose the distribution P of $(\boldsymbol{x}, y)$ is such that, conditioned on $\boldsymbol{x}$,

$$y|\boldsymbol{x} \sim \mathrm{N}\big(\langle \boldsymbol{x}, \boldsymbol{w}_\star \rangle, \sigma^2\big).$$

- **Question**: Given i.i.d. sample S, what is the MLE for $\boldsymbol{w}_\star$?

Maximum likelihood interpretation

Suppose the distribution P of $(\boldsymbol{x}, y)$ is such that, conditioned on $\boldsymbol{x}$,

$$y|\boldsymbol{x} \sim \mathrm{N}\big(\langle \boldsymbol{x}, \boldsymbol{w}_\star \rangle, \sigma^2\big).$$

- **Question**: Given i.i.d. sample S, what is the MLE for $\boldsymbol{w}_\star$?
- **Answer**: The ordinary least squares estimator.

  Log-likelihood of $\boldsymbol{w}$ given S:

$$\sum_{(\boldsymbol{x},y)\in S} \ln\left\{\exp\left(-\frac{1}{2}\Big(y - \langle \boldsymbol{x}, \boldsymbol{w}\rangle\Big)^2\right)\right\}$$

  (plus terms that don't depend on $\boldsymbol{w}$).

  $\longrightarrow$ maximizing likelihood $\equiv$ minimizing empirical mean squared error.

Taking $\boldsymbol{w}_\star$ seriously

Sometimes people seriously interpret the estimated regression coefficients $\hat{\boldsymbol{w}}_{\text{ols}}$.

Taking $\boldsymbol{w}_\star$ seriously

Sometimes people seriously interpret the estimated regression coefficients $\hat{\boldsymbol{w}}_{\text{ols}}$.

Example:

$$\hat{w}_{\text{ols},i} = 0 \quad \longrightarrow \quad \text{variable } x_i \text{ has negligible effect on } y \text{ in } \boldsymbol{w}_\star.$$

$$|\hat{w}_{\text{ols},i}| \gg 0 \quad \longrightarrow \quad \text{variable } x_i \text{ has significant effect on } y \text{ in } \boldsymbol{w}_\star.$$

Taking $\boldsymbol{w}_\star$ seriously

Sometimes people seriously interpret the estimated regression coefficients $\hat{\boldsymbol{w}}_{\text{ols}}$.

Example:

$$\hat{w}_{\text{ols},i} = 0 \quad \longrightarrow \quad \text{variable } x_i \text{ has negligible effect on } y \text{ in } \boldsymbol{w}_\star.$$
$$|\hat{w}_{\text{ols},i}| \gg 0 \quad \longrightarrow \quad \text{variable } x_i \text{ has significant effect on } y \text{ in } \boldsymbol{w}_\star.$$

Hypothesis tests for this usually assume $y|\boldsymbol{x} \sim \mathrm{N}(\langle \boldsymbol{x}, \boldsymbol{w}_\star \rangle, \sigma^2)$.

Computation

Fast computation

Naïve computation takes $O(np^2)$ time.

Fast computation

Naïve computation takes $O(np^2)$ time.

Hopes for speeding things up:

Fast computation

Naïve computation takes $O(np^2)$ time.

Hopes for speeding things up:

- Exploit sparsity or other structure in $\boldsymbol{X}$.

Fast computation

Naïve computation takes $O(np^2)$ time.

Hopes for speeding things up:

- Exploit sparsity or other structure in $\boldsymbol{X}$.
- Use iterative methods (e.g., *conjugate gradient* method).

Recap

- **Ordinary least squares** (three perspectives):
  1. Approximates $\boldsymbol{y}$ as linear combination of columns of $\boldsymbol{X}$.
  2. In statistical learning, excess mean squared error is $O(p/n)$.
  3. Maximum likelihood estimator under a Gaussian assumption.

Recap

- **Ordinary least squares** (three perspectives):
  1. Approximates $\boldsymbol{y}$ as linear combination of columns of $\boldsymbol{X}$.
  2. In statistical learning, excess mean squared error is $O(p/n)$.
  3. Maximum likelihood estimator under a Gaussian assumption.
- What if $p > n$?

Recap

- **Ordinary least squares** (three perspectives):
  1. Approximates $\boldsymbol{y}$ as linear combination of columns of $\boldsymbol{X}$.
  2. In statistical learning, excess mean squared error is $O(p/n)$.
  3. Maximum likelihood estimator under a Gaussian assumption.
- What if $p > n$?

  Ordinary least squares no longer uniquely defined, and often ill-behaved.

Recap

- **Ordinary least squares** (three perspectives):
  1. Approximates $\boldsymbol{y}$ as linear combination of columns of $\boldsymbol{X}$.
  2. In statistical learning, excess mean squared error is $O(p/n)$.
  3. Maximum likelihood estimator under a Gaussian assumption.
- What if $p > n$?

  Ordinary least squares no longer uniquely defined, and often ill-behaved.

  Use regularization:
  - force $\|\boldsymbol{w}\|_2^2$ to be small ("ridge regression") $\rightarrow$ can kernelize this
  - force $\|\boldsymbol{w}\|_1$ to be small ("Lasso")
  - force $\boldsymbol{w}$ to be sparse ("sparse regression")

Recap

- **Ordinary least squares** (three perspectives):
  1. Approximates $\boldsymbol{y}$ as linear combination of columns of $\boldsymbol{X}$.
  2. In statistical learning, excess mean squared error is $O(p/n)$.
  3. Maximum likelihood estimator under a Gaussian assumption.
- What if $p > n$?

  Ordinary least squares no longer uniquely defined, and often ill-behaved.

  Use regularization:
  - force $\|\boldsymbol{w}\|_2^2$ to be small ("ridge regression") $\rightarrow$ can kernelize this
  - force $\|\boldsymbol{w}\|_1$ to be small ("Lasso")
  - force $\boldsymbol{w}$ to be sparse ("sparse regression")

  We'll study OLS, RR, Lasso in a simplified "fixed-design" setting.