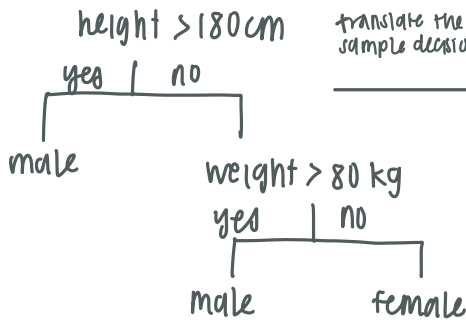

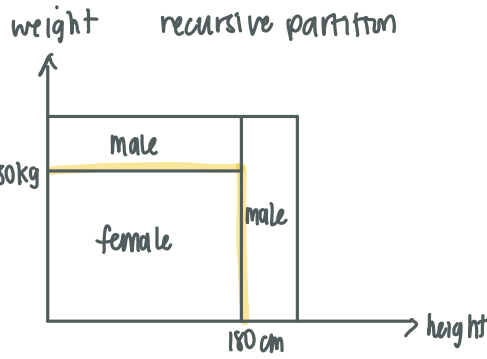


Trees


classification decision tree



translate the sample decision tree



\hat{p} = prob(male) for a region
 gini: $\hat{p}(1-\hat{p})$
 (x samples in the region)



each partition leads to maximal improvement in overall purity

gini index: weight each partitioning by n_{region} to observe reduction

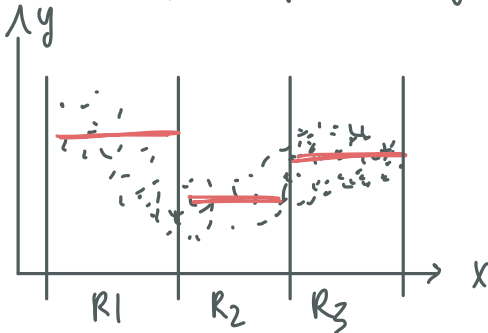
each iteration:

- choose a region
- choose a variable to partition on
- choose a cut point

exhaustive search to determine optimal partitioning for recursive partition
 a heuristic search

Regression trees

1-D example, single x & single y



max reduction in $Loss(R_1, \dots, R_m, \dots, R_p)$

stop if reduction $< \lambda$.


$$y = f(x) + e$$

$$f(x) = c_1 \mathbb{1}(x \in R_1) + c_2 \mathbb{1}(x \in R_2) + c_3 \mathbb{1}(x \in R_3)$$

$$\mathbb{1} = \begin{cases} 1 & \text{if } x \in R_1 \\ 0 & \text{if } x \notin R_1 \end{cases}$$

Regression trees (continued)

x_2 2-D example



Least squares loss

$$Loss = \sum_{i=1}^n (y_i - f(x_i))^2$$

$$= \sum_{m=1}^M \sum_{x_i \in R_m} (y_i - c_m)^2 + \lambda M$$

$$Loss(R_1, \dots, R_m, \dots, R_M)$$

piecewise approximation

↑
add penalty to grow shadow tree

$$\frac{d}{dc_m} = \sum_{x_i \in R_m} (y_i - c_m) = 0$$


$$\sum_{x_i \in R_m} y_i - N_m$$

↓
number of $x_i \in R_m$

$$\hat{c}_m = \frac{\sum_{x_i \in R_m} y_i}{N_m} = \frac{\sum_{i=1}^n y_i \mathbb{1}(x_i \in R_m)}{\sum_{i=1}^n \mathbb{1}(x_i \in R_m)}$$

Forest



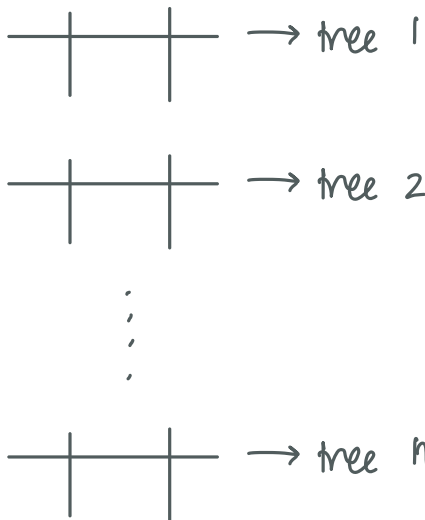
each split
go over a subset of variables



Random sampling

with Replacement

↓
bootstrap




boosting

ensemble machine
committee machine
add trees sequentially

$$f(x) = \sum_{k=1}^K h_k(x)$$

$h_k(x)$ is a tree



$$\text{Loss} = \sum_{i=1}^n (y_i - f(x_i))^2$$

current committee

$$s_i = \underbrace{\sum_{k=1}^{t-1} h_k(x_i)}_{\hat{s}_i} + h_t(x_i)$$

current committee (fixed) new member (tree) to be grown

$$\begin{aligned} \text{Loss} &= \sum_{i=1}^n (y_i - s_i)^2 \\ &= \sum_{i=1}^n (y_i - (\hat{s}_i + h_t(x_i)))^2 \\ &= \sum_{i=1}^n (\underbrace{y_i - \hat{s}_i}_{\hat{e}_i} - h_t(x_i))^2 \\ &= \sum_{i=1}^n (\hat{e}_i - h_t(x_i))^2 \end{aligned}$$

treat \hat{e}_i as response

small jumps:

$$\text{Loss} = \sum_{m=1}^M \sum_{x_i \in R_m} (\hat{e}_i - c_m)^2 + \lambda M + \sum_{m=1}^M \gamma c_m^2$$

we penalize c_m

gamma

Logistic regression

Iterated reweighted least squares

$$\sum_{i=1}^n \hat{w}_i (\hat{y}_i - h_t(x_i))^2$$

\uparrow $\leftarrow [x_i^T \Delta \beta \text{ substitute}]$
 $\frac{\hat{e}_i}{\hat{w}_i} \leftarrow y_i - \hat{p}_i \leftarrow \frac{e^{\hat{s}_i}}{1 + e^{\hat{s}_i}}$
 $\hat{w}_i \leftarrow \hat{p}_i (1 - \hat{p}_i)$

extreme gradient boosting

Instead of finding $\Delta \beta$, we grow a new tree to expand the error

each tree outputs a continuous score & does not withhold info compared to classification trees (binary)

Adaboost

less commonly used than extreme gradient boosting

$$s_i = f(x_i) = \sum_{k=1}^K \beta_k \boxed{h_k(x_i)}$$

weak classifier
 \downarrow
 classification tree $\rightarrow \{+1, -1\}$

each step requires to grow a new tree + calculate $\hat{\beta}$

$$\hat{y}_i = \text{sign}(s_i)$$