Lecture 12.

- Basic Machine Learning
  - Optimization
  - Convexity
  - Gradient Descent

Roadmap

① Privacy Attacks
- Reconstruction attacks
- Attack on K-Anon.
  (Composition).

② Differential Privacy.
- Randomized Response
- Laplace Mech.
- Exp. Mech.
- Gaussian Mech.

Composition

Post-Processing

Group Privacy

③ Applications.
- DP ML.
- DP Synthetic Data.

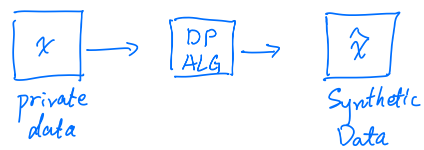

④ Fairness in ML.
- Consequential decision-making

⑤ Cryptography. (?)
Local Model of DP.

Linear Regression

	total_bill	tip
0	16.99	1.01
1	10.34	1.66
2	21.01	3.50
3	23.68	3.31
4	24.59	3.61
5	25.29	4.71
6	8.77	2.00
7	26.88	3.12
8	15.04	1.96
9	14.78	3.23
10	10.27	1.71
11	35.26	5.00
12	15.42	1.57
13	18.43	3.00
14	14.83	3.02
15	21.58	3.92

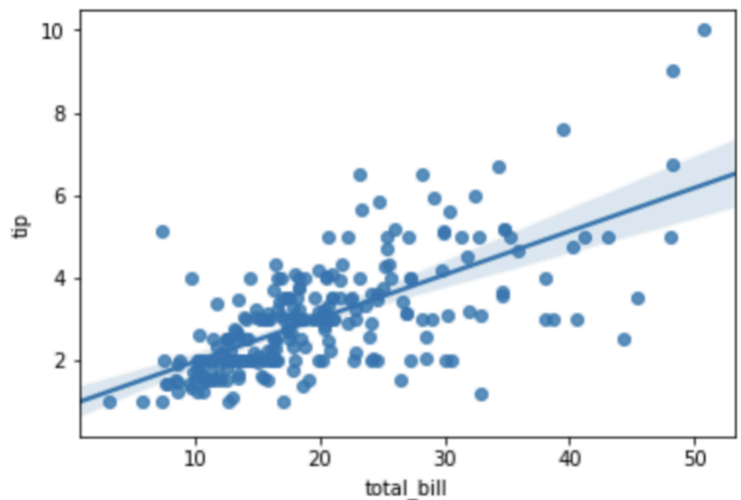

labeled example (x_i, y_i)

x_i = total bill

y_i = tip.

$$\min_{w \in \mathbb{R}} \sum_{i=1}^{n} \underbrace{(w x_i - y_i)^2}_{\text{squared loss.}}$$

Linear Classification

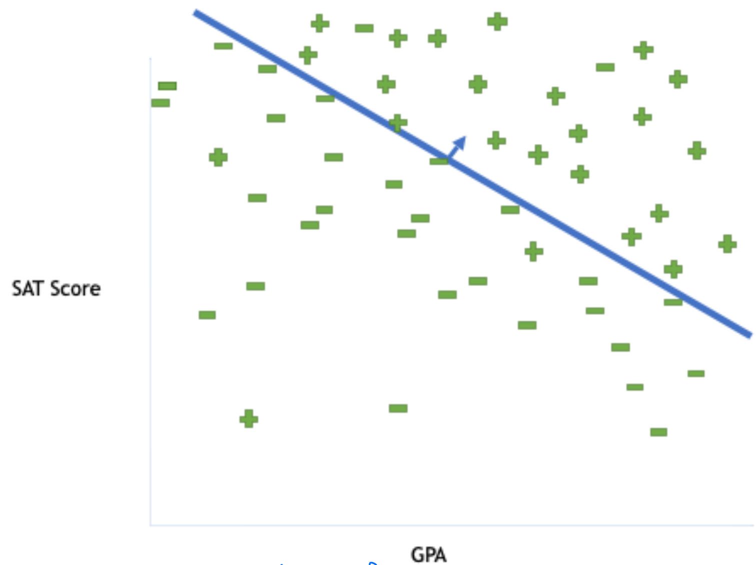

Candidates for loss functions:

0-1 loss : $\sum_{i=1}^{n} \mathbb{1}[\ \text{sign}(\langle w, x_i \rangle) \neq y_i]$ ← NP-hard to optimize.

indicator. feature (SAT, GPA) label. {1,-1}

$\langle a, b \rangle = a_1 b_1 + a_2 b_2$ (in 2-d).

$z_i = y_i \langle w, x_i \rangle$

± 1 want to match sign of y_i

Logistic Loss : $\sum_{i=1}^{n} \ln(1 + \exp(-z_i))$.

(Private) Optimization.

Given a data set $x = (x_1, \dots, x_n)$

loss function: ℓ

feasible set of parameters (weights): $C \subseteq \mathbb{R}^d$

Empirical Risk Minimization (ERM):

Lambda!

$$\min_{w \in C} \underbrace{L(w; x)}_{\text{Empirical Risk.}} = \frac{1}{n} \sum_{i=1}^{n} \ell(w; x_i) + \underbrace{\Lambda(w)}_{\text{Optional Regularization}}$$

for example: $\lambda \|w\|_2^2$

Goal: Find $\hat{w} \in C$ such that

$$L(\hat{w}, x) - \min_{w \in C} L(w, x)$$

is "small".

"Regret"

difference w.r.t. the "best"

Empirical Risk: $L(w; x) = \frac{1}{n} \sum_{i=1}^{n} \ell(w; x_i)$

Population Risk: $L(w; P) = \mathbb{E}_{x' \sim P} [\ell(w; x')]$

Usually, ER is a good proxy for PR.

Differential Privacy $\Longrightarrow$ Preventing overfitting.

Examples of loss functions:

$x = ((x_1, y_1), \ldots, (x_n, y_n))$

Squared loss $L(w; x) = \frac{1}{n} \sum_{i=1}^{n} (\langle w, x_i \rangle - y_i)^2$

Hinge Loss $L(w; x) = \frac{1}{n} \sum_{i=1}^{n} (1 - y_i \langle w, x_i \rangle)_+$

(used in support vector machine)

$$(a)_+ = \begin{cases} a & \text{if } a > 0 \\ 0 & \text{o/w.} \end{cases}$$

2 Criteria: ① Captures predictive accuracy

② Easy to optimize.

Convexity. (Sets and functions)

- $C \subseteq \mathbb{R}^d$ is convex if $\forall x, y \in C$, $t \in [0,1]$

$$\underbrace{t \cdot x + (1-t) \cdot y}_{\text{line segment}} \in C$$

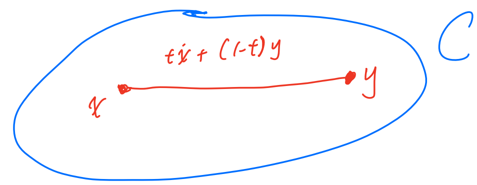

- $f : C \longrightarrow \mathbb{R}$ is convex if $\forall x, y \in C$

$$\underbrace{f\left(\frac{x+y}{2}\right)} \leq \underbrace{\frac{f(x)+f(y)}{2}}$$

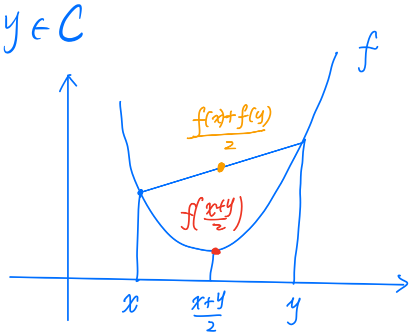

"Local Search"

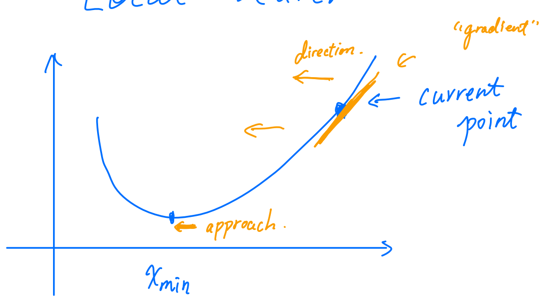

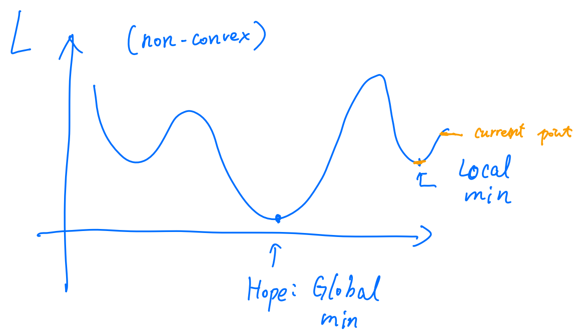

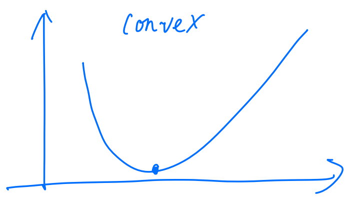

$\ell : C \times X \longrightarrow \mathbb{R}$

$\ell(w, x)$ measures "loss"

$L : C \longrightarrow \mathbb{R}$

$$L(w) = \frac{1}{n} \sum_{i=1}^{n} \ell(w, x_i)$$

$$\Pi_C(w) = \underset{w' \in C}{\arg\min} \; \|w - w'\|_2$$

"Projection"