Georgia Tech

Machine Learning CS 4641

Probability and Statistics

Nimisha Roy

Lecturer, SCI, College of Computing, Georgia Tech

Director, Online Undergraduate Initiatives

Slides adapted from Max Roozbahani, Le Song , Sam Roweis, and Chao Zhang.

Georgia Tech

Outline

- Probability Distributions
- Joint and Conditional Probability Distributions
- Bayes' Rule
- Mean and Variance
- Properties of Gaussian Distribution
- Maximum Likelihood Estimation

Probability

- A sample space S is the set of all possible outcomes of a conceptual or physical, repeatable experiment. (S can be finite or infinite.)
  - E.g., S may be the set of all possible outcomes of a dice roll: S $(1 \quad 2 \quad 3 \quad 4 \quad 5 \quad 6)$
  - E.g., S may be the set of all possible nucleotides of a DNA site: S $(A \quad C \quad G \quad T)$

  - E.g., S may be the set of all possible time-space positions of an aircraft on a radar screen.
- An Event A is any subset of S
  - Seeing "1" or "6" in a dice roll; observing a "G" at a site; UA007 in space-time interval

Three Key Ingredients in Probability Theory

A **sample space** is a collection of all possible **outcomes**

Random variables X represents **outcomes** in sample space

random variable

Probability of a random variable to happen $p(x) = p(X = x)$

1 specific outcome

$$p(x) \geq 0$$

Continuous variable

probability distribution function

Definition: Takes values from a continuous range (e.g., any real number within an interval).

Distribution: Governed by a **probability density function** (PDF).

Key Concepts:

- The **density** represents likelihood, but not actual probability at a specific point.
- Example: **Temperature**, which can be any real value (e.g., 72.3°F).
- Common distribution: **Gaussian (Normal) Distribution**.

$$\int_x p(x)dx = 1$$

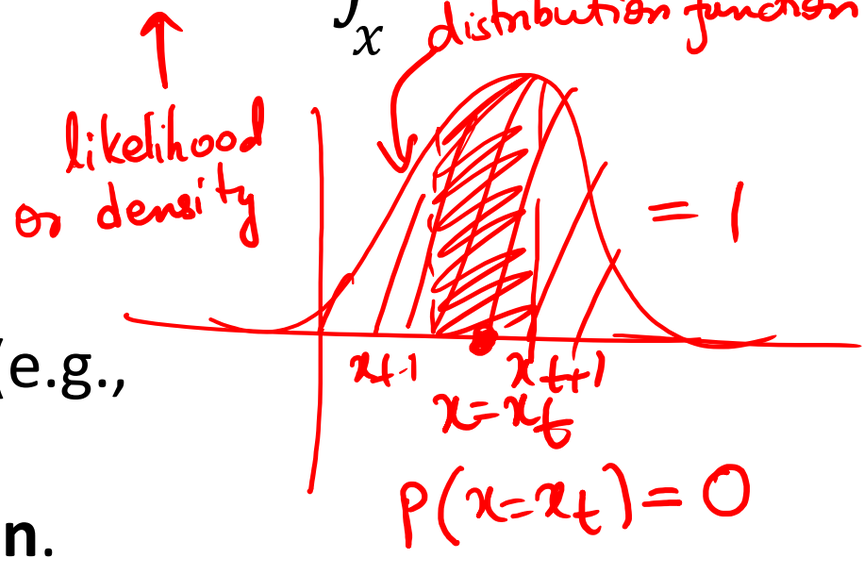

Discrete variable

Definition: Takes values from a countable set (e.g., integers).

Distribution: Governed by a **probability mass function** (PMF).

Key Concepts:

- The function directly gives **probability values**.
- Example: **A coin flip** (e.g., 0 for tails, 1 for heads).
- Common distribution: **Bernoulli Distribution**.

$$\sum_{x \in A} p(x) = 1$$

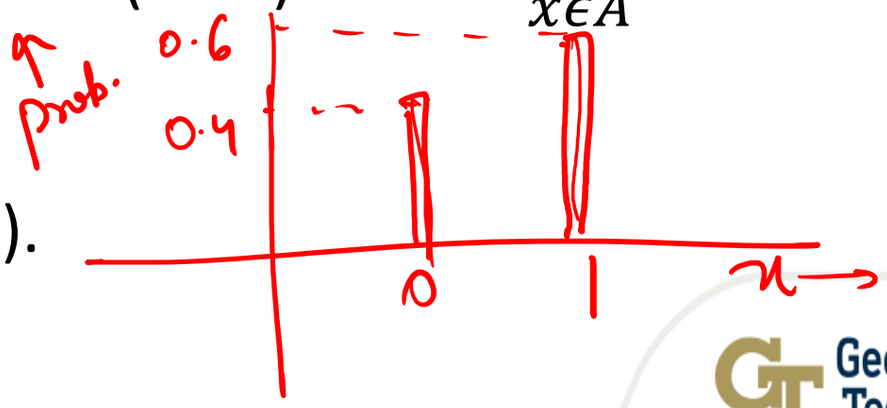

Georgia Tech

$f(x) \rightarrow$ objective function

$\theta \rightarrow$ parameters wt.

Continuous Probability Functions

$\frac{1}{7}(8-1) = 1$

- Examples:
  - Uniform Density Function:

$$f_x(x) = \begin{cases} \frac{1}{b-a} & for\ a \le x \le b \\ 0 & otherwise \end{cases}$$

$\theta \in (a, b)$ (min, max)

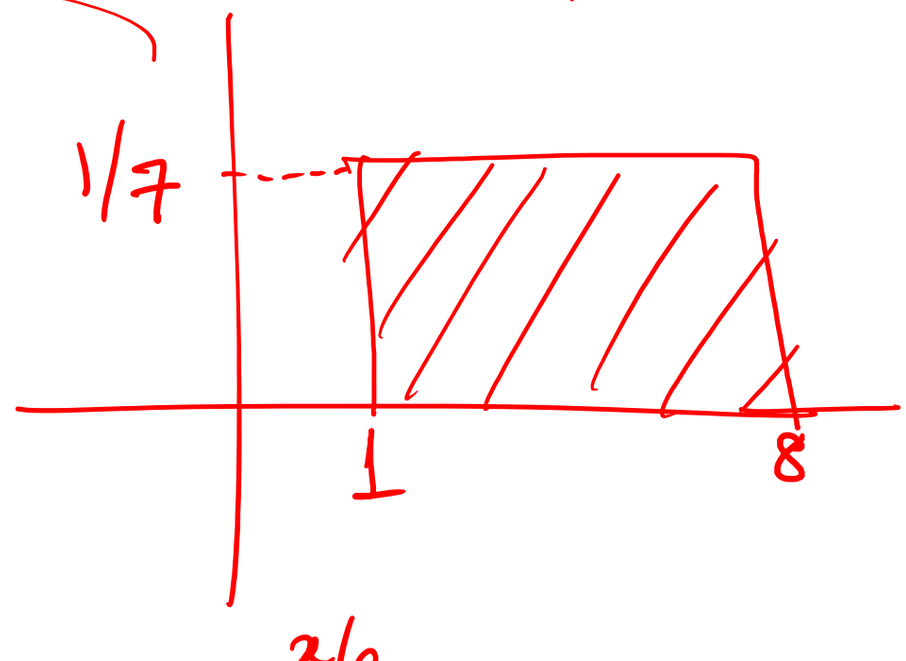

  - Exponential Density Function:

$$f_x(x) = \frac{1}{\mu} e^{-\frac{x}{\mu}} \qquad for\ x \ge 0$$

$\theta \in a$

$f(x) = \frac{1}{a} e^{-x/a}$

$a = \mu$

  - Gaussian(Normal) Density Function

$$f_x(x) = \frac{1}{\sqrt{2\pi}\sigma} e^{-\frac{(x-\mu)^2}{2\sigma^2}}$$

Discrete Probability Functions

- Examples:
  - Bernoulli Distribution:
    - $\begin{cases} 1-p & for\ x=0 \\ p & for\ x=1 \end{cases}$

In Bernoulli, just a **single** trial is conducted

  - Binomial Distribution:
    - $P(X=k)=\binom{n}{k}p^k(1-p)^{n-k}$

k is number of successes

n-k is number of failures

$\binom{n}{k}$ The total number of ways of selection **k** distinct combinations of **n** trials, **irrespective of order**.

E.g. - Given a biased coin with $\mu = 0.3$ for tail $(x = 1)$, what is the probability of getting 4 heads given we flip the coin 10 times?

Georgia Tech

Outline

- Probability Distributions
- Joint and Conditional Probability Distributions
- Bayes' Rule
- Mean and Variance
- Properties of Gaussian Distribution
- Maximum Likelihood Estimation

Example

X = Throw a dice

Y = Flip a coin

X and **Y** are random variables

N = total number of trials

n_{ij} = Number of occurrence

X

Y	$x_{i=1}=1$	$x_{i=2}=2$	$x_{i=3}=3$	$x_{i=4}=4$	$x_{i=5}=5$	$x_{i=6}=6$	C_j
$y_{j=2}=tail$	$n_{ij}=3$	$n_{ij}=4$	$n_{ij}=2$	$n_{ij}=5$	$n_{ij}=1$	$n_{ij}=5$	20
$y_{j=1}=head$	$n_{ij}=2$	$n_{ij}=2$	$n_{ij}=4$	$n_{ij}=2$	$n_{ij}=4$	$n_{ij}=1$	15
C_i	5	6	6	7	5	6	N=35

X

Y	$x_{i=1}=1$	$x_{i=2}=2$	$x_{i=3}=3$	$x_{i=4}=4$	$x_{i=5}=5$	$x_{i=6}=6$	C_j
$y_{j=2}=tail$	$n_{ij}=3$	$n_{ij}=4$	$n_{ij}=2$	$n_{ij}=5$	$n_{ij}=1$	$n_{ij}=5$	20
$y_{j=1}=head$	$n_{ij}=2$	$n_{ij}=2$	$n_{ij}=4$	$n_{ij}=2$	$n_{ij}=4$	$n_{ij}=1$	15
C_i	5	6	6	7	5	6	N=35

Joint probabilities $P(x=2, y=t) = 4/35 = n_{ij}/N$

Marginal prob. $\begin{cases} P(x=2) = 6/35 = C_i/N \\ P(y=t) = 20/35 = C_j/N \end{cases}$

sum rule $\Rightarrow$ $p(x) = \sum_y p(x,y)$

$p(y) = \sum_x p(x,y)$

$p(y) = \sum_{x,z,w} p(x,y,z,w)$

Cond. prob. $\begin{cases} P(x=2 \mid y=t) = 4/20 = \frac{n_{ij}}{C_j} \\ P(y=t \mid x=2) = 4/6 = \frac{n_{ij}}{C_i} \end{cases}$ Shrunk sample space

product rule = $P(x,y) = P(x|y)\cdot P(y)$

$= P(y|x)\cdot P(x)$

$P(x,y,z) = P(x|y,z)P(y,z)$

Georgia Tech

Probability:

$$p(X = x_i) = \frac{c_i}{N}$$

Joint probability:

$$p(X = x_i, Y = y_j) = \frac{n_{ij}}{N}$$

Conditional probability:

$$p(Y = y_j | X = x_i) = \frac{n_{ij}}{c_i}$$

Sum rule

$$p(X = x_i) = \sum_{j=1}^{L} p(X = x_i, Y = y_j) \Rightarrow \quad p(X) = \sum_{Y} P(X, Y)$$

Product rule

$$p(X = x_i, Y = y_j) = \frac{n_{ij}}{N} = \frac{n_{ij}}{c_i}\frac{c_i}{N} = p(Y = y_j | X = x_i)p(X = x_i)$$

$$p(X, Y) = p(Y|X)p(X)$$

Georgia Tech

If $P(x|y) = P(y) \Rightarrow$ x is conditionally ind. of y

Conditional Independence

- Examples:

$P(h, f, d, v)$

$= P(h|f,v,d)\,P(f,v|d)$

$= P(h|f,d)\,P(f|v,d)\,P(v|d)$

$= P(h|f,d)\,P(f|v)\,P(v|d)$

$P(Virus|DrinkBeer) = P(Virus)$

iff Virus is independent of Drink Beer

$P(Flu|Virus, DrinkBeer) = P(Flu|Virus)$

iff Flu is independent of Drink Beer, given Virus

$P(Headache|Flu, Virus, DrinkBeer) = P(Headache|Flu, DrinkBeer)$

$= P(h|f,d)\,P(f|v)\,P(v|d)\,P(d)$

$\leftarrow P(h|f,d)\,P(f|v)\,P(v)\,P(d)$

iff Headache is independent of Virus, given Flu and Drink Beer

Assume the above independence, we obtain:

$P(Headache, Flue, Virus, DrinkBeer)$

$= P(Headache|Flu, Virus, DrinkBeer)\,P(Flu|Virus, DrinkBeer)$

$P(Virus|DrinkBeer)P(DrinkBeer)$

$= P(Headache|Flu, DrinkBeer)P(Flu|Virus)P(Virus)P(DrinkBeer)$

Recap

- Sample Space and Event

all possible outcomes

1 possible outcome

- Probability Distribution Function

continous → prob. density function

discrete → prob. mass function

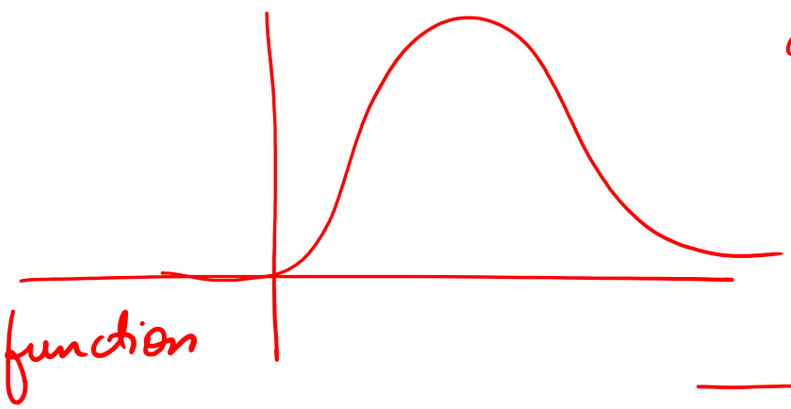

area under the curve is 1

OR

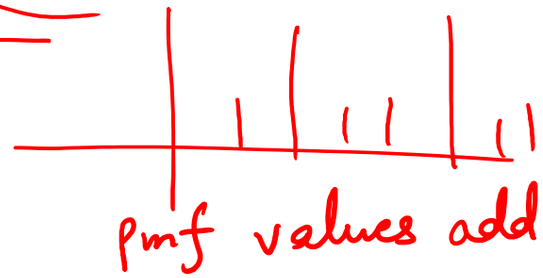

pmf values add to 1

- Objective Function

$f(x)$ ← optimize it → θ parameters.

- Joint Probability, Marginal Probability & Conditional Probability

$p(x, y, z, \ldots)$ — $p(x)$ — $p(x \mid y, z, \ldots)$

- 2 rules

→ SUM RULE →

$$p(x) = \sum_{y} p(x, y)$$

$$p(x) = \sum_{y,z,w} p(x, y, z, w)$$

→ PRODUCT RULE

$$p(x, y) = p(x \mid y)\, p(y)$$

$$= p(y \mid x)\, p(x)$$

$p(x \mid y) = p(x)$

x & y are conditionally independence

Georgia Tech

Outline

- Probability Distributions
- Joint and Conditional Probability Distributions
- Bayes' Rule
- Mean and Variance
- Properties of Gaussian Distribution
- Maximum Likelihood Estimation

Mix of sum & product rule

Bayes' Rule

$$P(x,y) = P(x|y)\,P(y) = P(y|x)\,P(x)$$

- P(X|Y)= Fraction of the worlds in which X is true given that Y is also true.

$$P(x|y) = \frac{P(x,y)}{P(y)}$$

$$P(x|y) = \frac{P(y|x)\,P(x)}{P(y)}$$

- For example:
  - H="Having a headache"
  - F="Coming down with flu"
  - $P(Headche|Flu)$ = fraction of flu-inflicted worlds in which you have a headache. How to calculate?

- Definition:

$$P(X|Y) = \frac{P(X,Y)}{P(Y)} = \frac{\mathrm{P(Y|X)P(X)}}{\mathrm{P(Y)}}$$

Corollary:

$$P(X,Y) = P(Y|X)P(X)$$

This is called Bayes Rule

Bayes' Rule

$$p(y|x) = \frac{p(x|y)\,p(y)}{p(x)}$$

$$p(x,z) = \sum_y p(x,z,y) = \sum_y p(x|z,y)\,p(z,y) = p(x|z,y)p(z,y) + p(x|z,\bar{y})\,p(z,\bar{y})$$

$$= \frac{p(x|y)\,p(y)}{\sum_y p(x,y)}$$

binary

- $P(Headache|Flu) = \frac{P(Headache,Flu)}{P(Flu)}$

$$= \frac{P(Flu|Headache)P(Headache)}{P(Flu)}$$

$$= \frac{p(x|y)\,p(y)}{\sum_y p(x|y)\,p(y)} = \frac{p(x|y)\,p(y)}{p(x|y)\,p(y) + p(x|\bar{y})\,p(\bar{y})}$$

Other cases:

- $P(Y|X) = \frac{P(X|Y)P(Y)}{P(X|Y)P(Y)+P(X|\neg Y)P(\neg Y)}$
- $P(Y = y_i|X) = \frac{P(X|Y)P(Y)}{\sum_{i\in S} P(X|Y = y_i)P(Y=y_i)}$
- $P(Y|X,Z) = \frac{P(X|Y,Z)P(Y,Z)}{P(X,Z)} = \frac{P(X|Y,Z)P(Y,Z)}{P(X|Y,Z)P(Y,Z)+P(X|\neg Y,Z)P(\neg Y,Z)}$

binary

$$p(y|x,z)\,p(x,z) = p(y,x,z)$$

$$p(y|x,z) = \frac{p(x,y,z)}{p(x,z)}$$

$$p(x,z) = \sum_y p(x,z,y)$$

binary

Outline

- Probability Distributions
- Joint and Conditional Probability Distributions
- Bayes' Rule
- Mean and Variance
- Properties of Gaussian Distribution
- Maximum Likelihood Estimation

Mean and Variance

$$E[g(x)] = \int_x g(x)\,p(x)\,dx$$

$$= \sum_{x_i} g(x_i)\,p(x_i)$$

- Expectation: The mean value, center of mass, first moment:

$$E_X[g(X)] = \int_{-\infty}^{\infty} g(x)p_X(x)dx = \mu$$

In ML, Expectation = Arithmetic Mean

- N-th moment: $g(x) = x^n$
- N-th central moment: $g(x) = (x-\mu)^n$
- Mean: $E_X[X] = \int_{-\infty}^{\infty} xp_X(x)dx$
  - $E[\alpha X] = \alpha E[X]$
  - $E[\alpha + X] = \alpha + E[X]$
- Variance(Second central moment): $Var(x) = E_X[(X - E_X[X])^2] = E_X[X^2] - E_X[X]^2$
  - $Var(\alpha X) = \alpha^2 Var(X)$
  - $Var(\alpha + X) = Var(X)$

$$Var = \frac{1}{n}\sum_i (x_i - \mu)^2$$

$$E\left[X - E[X]\right]^2$$

Expected value and average

Why arithmetic mean is same as expected value:

3 sided biased die: $X = [1, 2, 3]$ $\quad p(x) = [1/6, 1/3, 1/2]$

valid pmf??

$g(x) = x$

$$E[g(x)] = \sum_i g(x_i) \cdot p(x_i)$$

$$= 1 \times \frac{1}{6} + 2 \times \frac{1}{3} + 3 \times \frac{1}{2} = \frac{1+4+9}{6} = \frac{14}{6}$$

$$\mu(x) = \frac{1+2+3}{3} = 2$$

Roll die six times → [1, 2, 2, 3, 3, 3]

$E[g(x)] = 14/6$

$\mu(x) = \frac{1+2+2+3+3+3}{6} = 14/6$

Expectation determined by probability distribution. Arithmetic average determined by observed outcomes of trials.

Both are the same in our class because we assume we have sufficient data that models the distribution. So, the arithmetic average is a good estimate of the true expectation (Law of Large Numbers).

Georgia Tech

Variance and average:

$$X = \begin{bmatrix} 1 \\ 2 \\ 3 \end{bmatrix}_{3 \times 1}^{n \times d}$$

$$Var = \frac{1}{N} \bar{X}^T \bar{X}$$

$$\bar{X} = \begin{bmatrix} 1-\mu_h \\ 2-\mu_h \\ 3-\mu_h \end{bmatrix}_{3 \times 1}$$

$$\frac{1}{N} \bar{X}^T \bar{X} = \frac{1}{N} \begin{bmatrix} 1-\mu_h & 2-\mu_h & 3-\mu_h \end{bmatrix} \begin{bmatrix} 1-\mu_h \\ 2-\mu_h \\ 3-\mu_h \end{bmatrix}$$

$$Var = \frac{1}{N} \left((1-\mu_h)^2 + (2-\mu_h)^2 + (3-\mu_h)^2 \right)$$

$$= \boxed{1 \times 1}$$

$$Var = \frac{1}{n} \sum_i (x_i - \mu)^2$$

$$E\left[x - E[x] \right]^2$$

Covariance:

$$X = \begin{bmatrix} h & w \\ 1 & 4 \\ 2 & 5 \\ 3 & 6 \end{bmatrix}$$

$Cov_{d \times d}$ for X with d features

$$\bar{X} = \begin{bmatrix} 1-\mu_h & 4-\mu_w \\ 2-\mu_h & 5-\mu_w \\ 3-\mu_h & 6-\mu_w \end{bmatrix}_{3 \times 2}$$

$$Cov = \frac{1}{N} \bar{X}^T \bar{X}$$

$$\frac{1}{N} \begin{bmatrix} 1-\mu_h & 2-\mu_h & 3-\mu_h \\ 4-\mu_w & 5-\mu_w & 6-\mu_w \end{bmatrix} \begin{bmatrix} 1-\mu_h & 4-\mu_w \\ 2-\mu_h & 5-\mu_w \\ 3-\mu_h & 6-\mu_w \end{bmatrix}$$

$$Cov(1,1) = \frac{1}{N}\left[(1-\mu_h)^2 + (2-\mu_h)^2 + (3-\mu_h)^2\right] = \sigma_h^2$$

$$Cov(2,2) = \frac{1}{N}\left[(4-\mu_w)^2 + (5-\mu_w)^2 + (6-\mu_w)^2\right] = \sigma_w^2$$

$$Cov(1,2) = \frac{1}{N}\left((1-\mu_h)(4-\mu_w) + (2-\mu_h)(5-\mu_w) + (3-\mu_h)(5-\mu_w)\right)$$

$$\sigma_{hw} = \sigma_{wh}$$

$Cov(2,1)$

If $\sigma_h^2 = \sigma_{hw} \Rightarrow h = w$

Georgia Tech

Covariance:

$$Cov = \begin{bmatrix} \sigma_h^2 & \sigma_{hw} \\ \sigma_{wh} & \sigma_{w^2} \end{bmatrix}$$

- Off-diagonal terms of covariance matrix same? Yes
- Diagonal terms higher in value or off diagonal term? ✓ Equal
- Will eigenvectors of covariance matrix be orthonormal? ✓
- Do we know its eigenvectors? V from SVD of X
- Any upper and lower bound in values of covariance matrix? No bound

Covariance measures how two features vary together. Problem: covariance depends on scale

Georgia Tech

Correlation:

$$\Sigma = \frac{1}{N} \bar{X}^{*T} \bar{X}^{*}$$

$X^* \rightarrow$ Standardized

Normalization $\rightarrow [0,1]$

Standardization $\rightarrow [-1,1]$

$$\frac{X - \mu}{\sigma}$$

$$X = \begin{bmatrix} 1 & 4 \\ 2 & 5 \\ 3 & 6 \end{bmatrix}$$

$$\bar{X}^{*} = \begin{bmatrix} \frac{1-\mu_h}{\sigma_h} & \frac{4-\mu_\omega}{\sigma_\omega} \\ \frac{2-\mu_h}{\sigma_h} & \frac{5-\mu_\omega}{\sigma_\omega} \\ \frac{3-\mu_h}{\sigma_h} & \frac{6-\mu_\omega}{\sigma_\omega} \end{bmatrix}$$

$$\Sigma = \frac{1}{N} \bar{X}^{*T} \bar{X}^{*}$$

$$\Sigma_{(1,1)} = \frac{1}{N} \frac{(1-\mu_h)^2 + (2-\mu_h)^2 + (3-\mu_h)^2}{\sigma_h^2}$$

$$= \frac{\sigma_h^2}{\sigma_h^2} = 1$$

$Cor(i,i) = 1$

$$\Sigma_{d \times d} = \begin{bmatrix} 1 & -1 < [\,] < 1 \\ & 1 \\ & & 1 \\ & & & 1 \end{bmatrix}$$

Exploratory Data Analysis

Useful in EDA. If features are correlated, then data may have redundancy.

Georgia Tech

For Joint Distributions

$$Cov(X,Y) = E[XY] - E[X]E[Y]$$

- Expectation and Covariance:
  - $E[X+Y] = E[X] + E[Y]$
  - $cov(X,Y) = E[(X - E_X[X])(Y - E_Y(Y)] = E[XY] - E[X]E[Y]$
  - $Var(X+Y) = Var(X) + 2cov(X,Y) + Var(Y)$

Uncorrelated vs Independent RV

Uncorrelated (Definition: Cov(X,Y)=0.)
- Means no **linear relationship** between X and Y.
- Does **not** rule out non-linear dependence.
- Example: X~U(−1,1), Y=x^2 → Uncorrelated but dependent.

Independent (Definition: P(X∈A,Y∈B)=P(X∈A)P(Y∈B))
- Stronger condition: knowing one variable gives **no information** about the other.
- Independence ⇒ Uncorrelated (if variances exist).

P(X, Y) = P(X) P(Y)

Key Differences
- Independence is a **stronger** property.
- Uncorrelated only removes **linear relationships**.
- Uncorrelated ⇏ Independent.

ML Relevance
- Most ML models care about uncorrelatedness because they only model linear relationships.
- True independence is rarer and much harder to check, but it's what we assume in stronger probabilistic models

Outline

- Probability Distributions
- Joint and Conditional Probability Distributions
- Bayes' Rule
- Mean and Variance
- Properties of Gaussian Distribution
- Maximum Likelihood Estimation

Gaussian Distribution

$$f(x|a,b) = \frac{1}{\sqrt{2\pi b}} e^{-\frac{(x-a)^2}{2b}}$$

$a = \mu$
$b = \sigma^2$

- Gaussian Distribution:

$$f(x|\mu,\sigma^2) = \frac{1}{\sqrt{2\pi\sigma^2}} e^{-\frac{(x-\mu)^2}{2\sigma^2}}$$

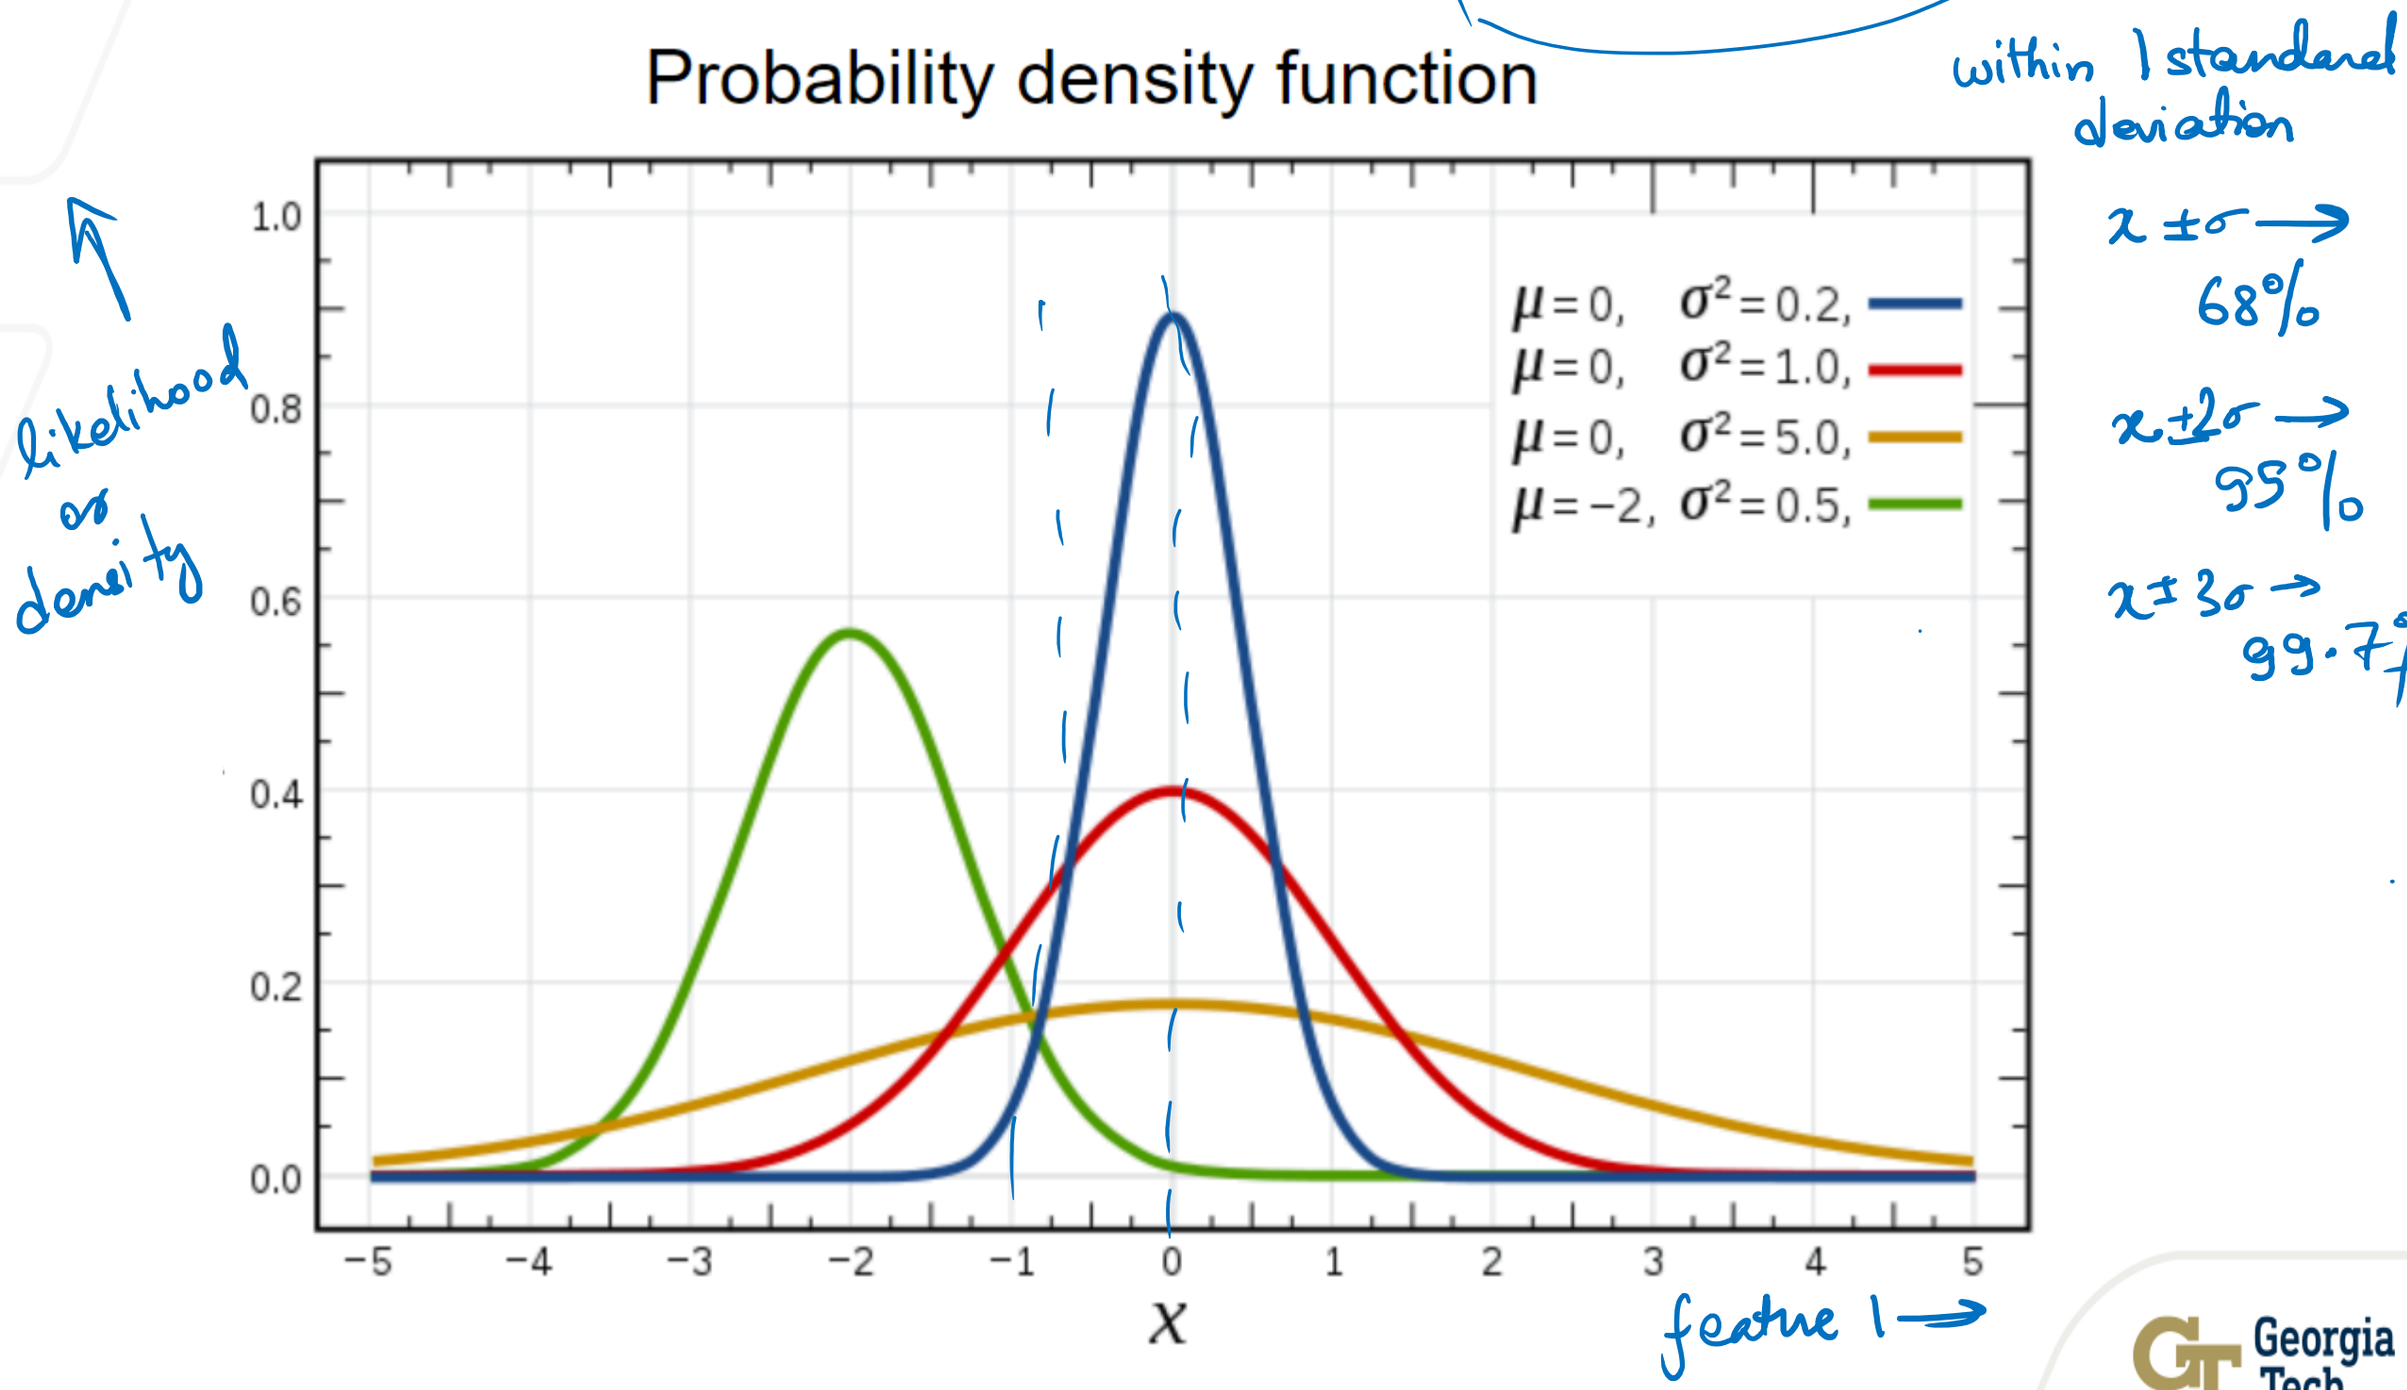

within 1 standard deviation

$x \pm \sigma \rightarrow$ 68%

$x \pm 2\sigma \rightarrow$ 95%

$x \pm 3\sigma \rightarrow$ 99.7%

Georgia Tech

Multivariate Gaussian Distribution

$$p(x|\mu,\Sigma) = \frac{1}{(2\pi)^{n/2}|\Sigma|^{1/2}}\exp\{-\frac{1}{2}(x-\mu)^{\top}\Sigma^{-1}(x-\mu)\}$$

- Moment Parameterization $\mu = E(X)$

$$\Sigma = Cov(X) = E[(X-\mu)(X-\mu)^{\top}]$$

- Mahalanobis Distance $\Delta^2 = (x-\mu)^{\top}\Sigma^{-1}(x-\mu)$

- Tons of applications (MoG, FA, PPCA, Kalman filter,…)

Properties of Gaussian Distribution

- The linear transform of a Gaussian r.v. is a Gaussian. Remember that no matter how x is distributed

$$E(AX + b) = AE(X) + b$$

$$Cov(AX + b) = ACov(X)A^{\top}$$

this means that for Gaussian distributed quantities:

$$X \sim N(\mu, \Sigma) \rightarrow AX + b \sim N(A\mu + b, A\Sigma A^{\top})$$

- The sum of two independent Gaussian r.v. is a Gaussian

$$Y = X_1 + X_2,\ X_1 \perp X_2 \rightarrow \mu_y = \mu_1 + \mu_2, \Sigma_y = \Sigma_1 + \Sigma_2$$

- The multiplication of two Gaussian functions is another Gaussian function (although no longer normalized)

$$N(a, A)N(b, B) \propto N(c, C),$$
$$where\ C = (A^{-1} + B^{-1})^{-1}, c = CA^{-1}a + CB^{-1}b$$

Central Limit Theorem

Probability mass function of a **biased** dice

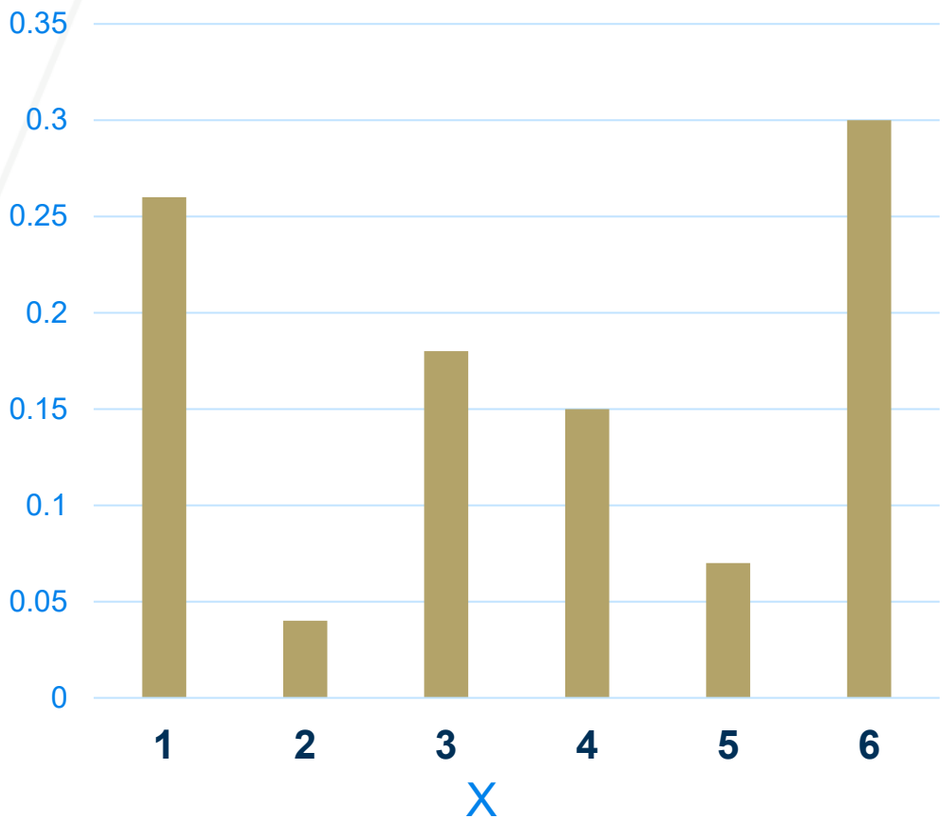

Let's say, I am going to get a sample from this pmf having a size of $\boldsymbol{n = 4}$

$$S_1 = \{1,1,1,6\} \Rightarrow E(S_1) = 2.25$$

$$S_2 = \{1,1,3,6\} \Rightarrow E(S_2) = 2.75$$

$$\vdots$$

$$S_m = \{1,4,6,6\} \Rightarrow E(S_m) = 4.25$$

- According to CLT, if you sample enough from any distribution with finite variance you will get an approximate Gaussian distribution.
- No matter what the population looks like, the average of many samples looks Normal.
- Explains **why the Normal distribution is everywhere** in statistics and ML.

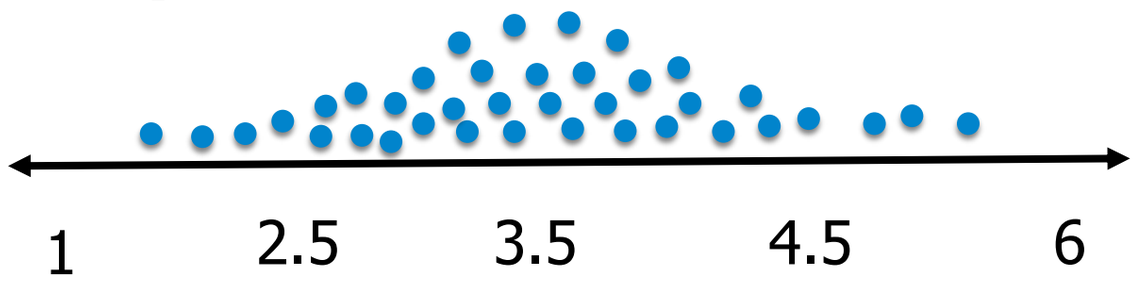

Georgia Tech

Prob vs Likelihood

Probability predicts data from parameters.
Likelihood evaluates parameters from data.

- **Probability**
  - Forward direction: given parameters, what's the chance of data?
  - $P(data \mid \theta)$
  - Example: *If the coin has bias* $\theta = 0.7$*, what's the probability of 8 heads in 10 tosses?*
  - Varies with **different outcomes** of data.
- **Likelihood**
  - Reverse view: given data, how plausible are parameters?
  - $L(\theta \mid data) \propto P(data \mid \theta)$
  - Example: *I observed 8 heads in 10 tosses. How likely is it that the coin's bias is* $\theta = 0.7$*?*
  - Varies with **different parameter values**.

Georgia Tech

Prob vs Likelihood

$\theta = \mu, \sigma^2$

Case A

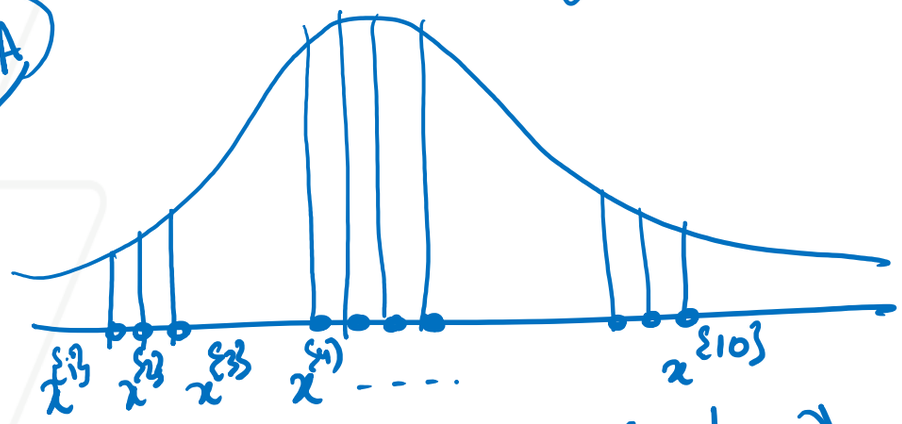

Case B

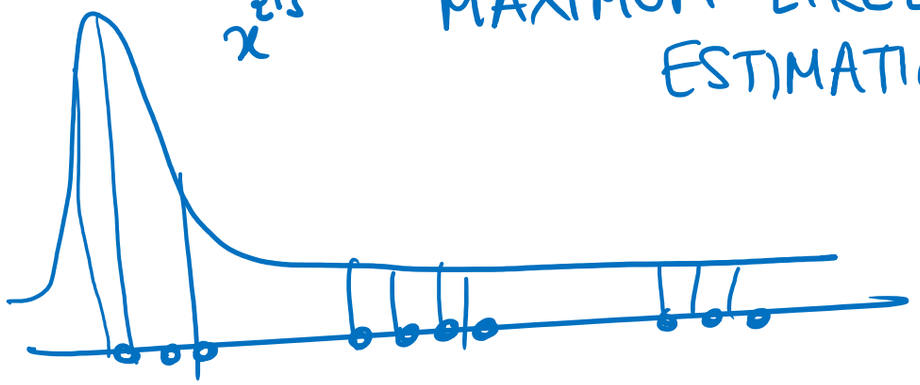

MAXIMUM LIKELIHOOD ESTIMATION

Likelihood → $L(\theta|x) = L(a, b|x)$

$= p(x^{\{1\}}, x^{\{2\}}, \ldots x^{\{10\}} | \theta)$

maximize

$L(\theta|x) = \max p(x^{\{1\}}, x^{\{2\}}, \ldots x^{\{10\}} | \theta)$

$= p(x^{\{1\}}|\theta) \cdot p(x^{\{2\}}|\theta) \ldots p(x^{\{10\}}|\theta)$

$= \prod_i p(x^{\{i\}}|\theta)$

iid assumption
↳ all datapoints are independent & identically distributed

DO NOT LIKE LIST

1. for loop
2. Inverse
3. Matrix. Mult.
4. Joint Prob.

Prob vs Likelihood

$\log abc = \log a + \log b + \log c$

$$L(\theta | x) = \prod_i p(x^{\{i\}} | \theta)$$

Maximize this probability

↓

most optimal values of a, b

$$\log L(\theta | x) = \log \prod_i p(x^{\{i\}} | \theta)$$

$$\log L(\theta | x) = \sum_i \log p(x^{\{i\}} | \theta)$$

$$\frac{\partial \log L}{\partial a} = 0 \qquad \frac{\partial \log L}{\partial b} = 0$$

a, b

Maximum Likelihood Estimation

Outline

- Probability Distributions
- Joint and Conditional Probability Distributions
- Bayes' Rule
- Mean and Variance
- Properties of Gaussian Distribution
- Maximum Likelihood Estimation

Maximum Likelihood Estimation

- Probability: inferring probabilistic quantities for data given fixed models (e.g. prob. of events, marginals, conditionals, etc).
- Statistics: inferring a model given fixed data observations (e.g. clustering, classification, regression).

Main assumption:

Independent and identically distributed random variables

i.i.d

$\theta = 0.5$

Maximum Likelihood Estimation

For Bernoulli (i.e. flip a coin):

Objective function: $P\left(x^{\{i\}}|\theta\right) = \theta^{x^{\{i\}}}(1-\theta)^{1-x^{\{i\}}}$ prob. of head $x^{\{i\}} \in \{0,1\}\ or\ \{head, tail\}$

$$L(\theta|X) = L(\theta|X = x^{\{1\}}, X = x^{\{2\}}, X = x^{\{3\}}, \dots, X = x^{\{n\}})$$

i.i.d assumption

$$L(\theta|X) = \prod_{i=1}^{n} P\left(x^{\{i\}}|\theta\right)$$

$$L(\theta|X) = \prod_{i=1}^{n} P\left(x^{\{i\}}|\theta\right) = \prod_{i=1}^{n} \theta^{x^{\{i\}}}(1-\theta)^{1-x^{\{i\}}}$$

$a^b \cdot a^c \cdot a^d = a^{b+c+d}$

$$L(\theta|X) = \theta^{x^{\{1\}}}(1-\theta)^{1-x^{\{1\}}} \times \theta^{x^{\{2\}}}(1-\theta)^{1-x^{\{2\}}} \dots \times \theta^{x^{\{n\}}}(1-\theta)^{1-x^{\{n\}}} =$$
$$= \theta^{\Sigma x^{\{i\}}}(1-\theta)^{\Sigma(1-x^{\{i\}})}$$

Georgia Tech

We don't like multiplication, let's convert it into summation

What's the trick? Take the log

$\log(a^b) = b \log a$

$$L(\theta|X) = \theta^{\sum x^{\{i\}}}(1-\theta)^{\sum(1-x^{\{i\}})}$$

$$logL(\theta|X) = l(\theta|X) = \log(\theta)\sum_{i=1}^{n} x^{\{i\}} + \log(1-\theta)\sum_{i=1}^{n}(1-x^{\{i\}})$$

How to optimize θ?

$$\frac{\partial l(\theta|X)}{\partial \theta} = 0 \qquad \frac{\sum_{i=1}^{n} x^{\{i\}}}{\theta} - \frac{\sum_{i=1}^{n} 1 - x^{\{i\}})}{1-\theta} = 0$$

$$\theta = \frac{1}{n}\sum_{i=1}^{n} x^{\{i\}}$$

- If your data has **70% ones**, the probability that best explains it is **θ = 0.7**
- Any other θ would make the observed pattern *less plausible*

Georgia Tech