

UVA CS 6316: Machine Learning

Lecture 17: Generative Bayes Classifiers

Dr. Yanjun Qi
University of Virginia
Department of Computer Science

Course Content Plan →

Six major sections of this course

~~Regression (supervised)~~

Y is a continuous

Classification (supervised)

Y is a discrete

Unsupervised models

NO Y

Learning theory

About $f()$

Graphical models

About interactions among X_1, \dots, X_p

Reinforcement Learning

Learn program to Interact with its environment

Three major sections for classification

- We can divide the large variety of classification approaches into **roughly three major types**

1. Discriminative

directly estimate a decision rule/boundary

e.g., ~~support vector machine~~, **decision tree**, ~~logistic regression~~,

e.g. ~~neural networks (NN), deep NN~~



2. Generative:

build a generative statistical model


e.g., Bayesian networks, **Naïve Bayes classifier**

3. ~~Instance based classifiers~~

- Use observation directly (no models)

- ~~e.g. K nearest neighbors~~

Today : Generative Bayes Classifiers



- ✓ Bayes Classifier (BC)
 - Generative Bayes Classifier
- ✓ Naïve Bayes Classifier
- ✓ Gaussian Bayes Classifiers
 - Gaussian distribution
 - Naïve Gaussian BC
 - Not-naïve Gaussian BC → LDA, QDA

Review: Bayes classifiers (BC)

- Treat each feature attribute and the class label as random variables.
- **Testing**: Given a sample \mathbf{x} with attributes (x_1, x_2, \dots, x_p) :
 - Goal is to predict its class c .
 - Specifically, we want to find the class that maximizes $p(c | x_1, x_2, \dots, x_p)$.
- **Training**: can we estimate $p(C_i | \mathbf{x}) = p(C_i | x_1, x_2, \dots, x_p)$ directly from data?

$$c_{MAP} = \operatorname{argmax}_{c_j \in C} P(c_j | x_1, x_2, \dots, x_p)$$

MAP Rule

Two kinds of Bayes classifiers via MAP classification rule

- Establishing a probabilistic model for classification
 - **(1) Discriminative**
 - **(2) Generative**

Review: Discriminative BC

Softmax
$$p(\hat{y}_i | \vec{x}) = \frac{e^{\beta_i}}{\sum_{j=1}^L e^{\beta_j}}$$

– (1) Discriminative model

$$\operatorname{argmax}_{c \in C} P(c | \mathbf{X}), \quad C = \{c_1, \dots, c_L\}$$

$$P(c_1 | \mathbf{x}) \quad P(c_2 | \mathbf{x}) \quad \dots \quad P(c_L | \mathbf{x})$$



$$x_1 \quad x_2 \quad \dots \quad x_p$$

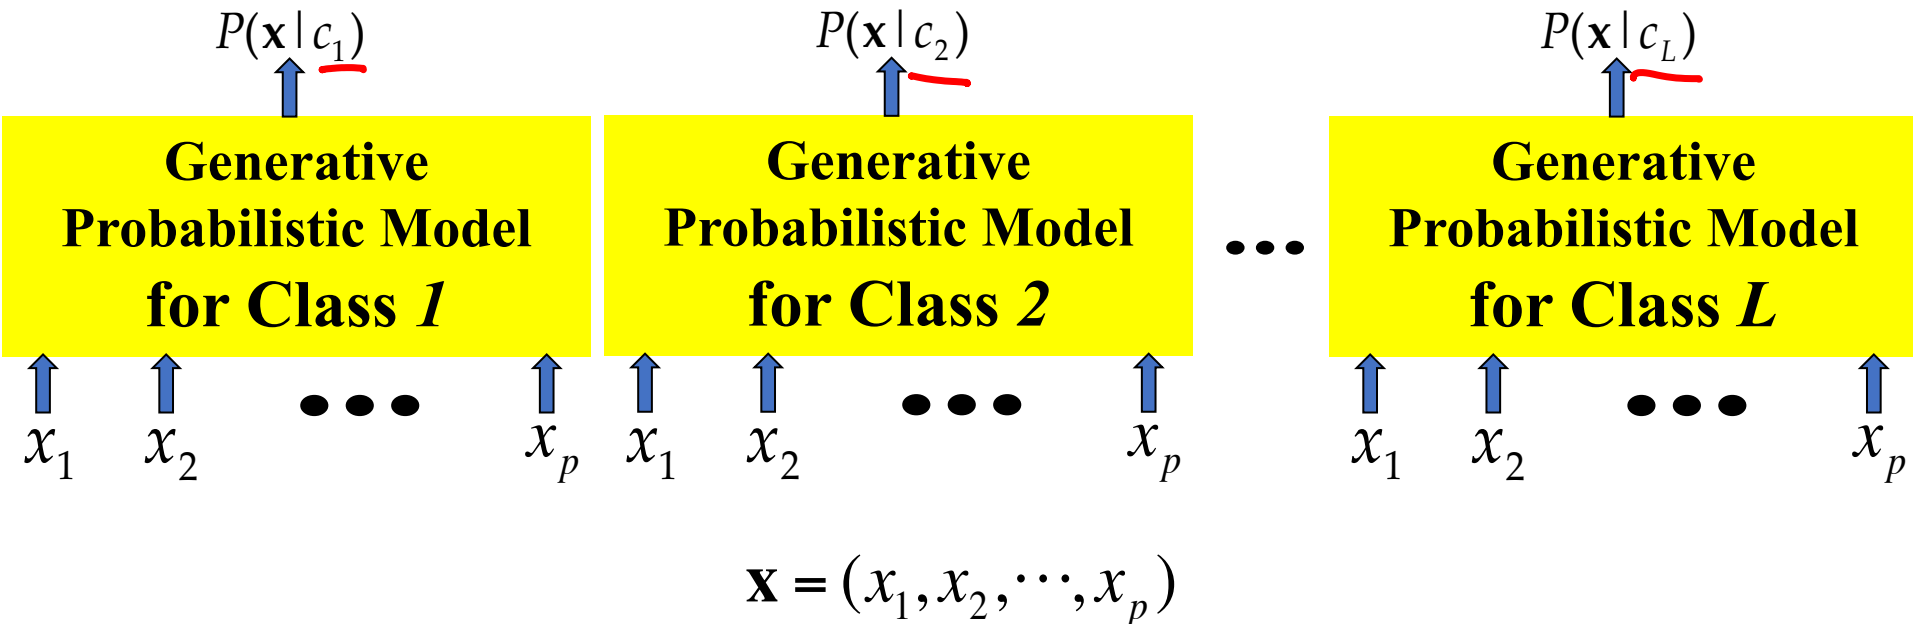
$$\mathbf{x} = (x_1, x_2, \dots, x_p)$$

multi-class
logistic
regression

(2) Generative BC

$$P(\mathbf{X} | C),$$
$$C = c_1, \dots, c_L, \mathbf{X} = (X_1, \dots, X_p)$$

$\Rightarrow p(c_i | \vec{x})$



Review Probability: If hard to directly estimate from data, most likely we can estimate

- 1. Joint probability
 - Use Chain Rule
- 2. Marginal probability
 - Use the total law of probability
- 3. Conditional probability
 - Use the Bayes Rule

$$p(C_i | \vec{x}) = \frac{p(\vec{x} | C_i) p(C_i)}{p(\vec{x})}$$

↓

$$C^* = \operatorname{argmax}_{C_i \in \{C_1, C_2, \dots, C_L\}} p(C_i | \vec{x})$$

Review : Bayes' Rule

– for Generative Bayes Classifiers

$$P(C | X) = \frac{P(X | C)P(C)}{P(X)}$$

$$P(C_1|x), P(C_2|x), \dots, P(C_L|x)$$

$$P(C_1), P(C_2), \dots, P(C_L)$$

$$P(C_i | \mathbf{X}) = \frac{P(\mathbf{X} | C_i)P(C_i)}{P(\mathbf{X})}$$

Review : Bayes' Rule

– for Generative Bayes Classifiers

$$P(C | X) = \frac{P(X | C)P(C)}{P(X)}$$

Prior

$P(C_1|x), P(C_2|x), \dots, P(C_L|x)$

$P(C_1), P(C_2), \dots, P(C_L)$

$$P(C_i | \mathbf{X}) = \frac{P(\mathbf{X} | C_i)P(C_i)}{P(\mathbf{X})}$$

Summary of Generative BC:

- Apply Bayes rule to get posterior probabilities

$$\operatorname{argmax}_{c_i} \left\{ P(C = c_i | \mathbf{X} = \mathbf{x}) \right\} = \frac{P(\mathbf{X} = \mathbf{x} | C = c_i) P(C = c_i)}{P(\mathbf{X} = \mathbf{x})}$$
$$\propto \frac{P(\mathbf{X} = \mathbf{x} | C = c_i) P(C = c_i)}{\text{for } i = 1, 2, \dots, L}$$

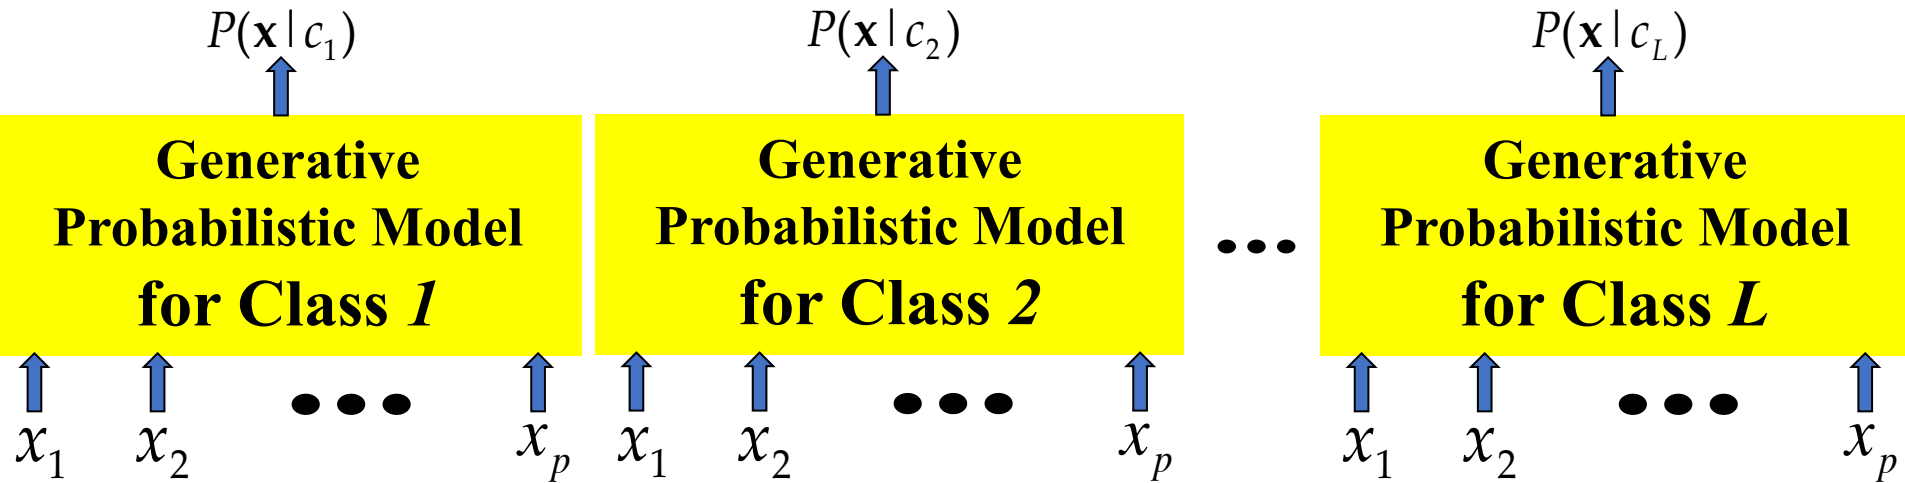
- Then apply the MAP rule

$$\left\{ \begin{array}{l} p(\vec{x} | c_i) \\ p(c_i) \end{array} \right. \quad c_i \in \{c_1, c_2, \dots, c_L\}$$

Establishing a probabilistic model for classification through generative probabilistic models

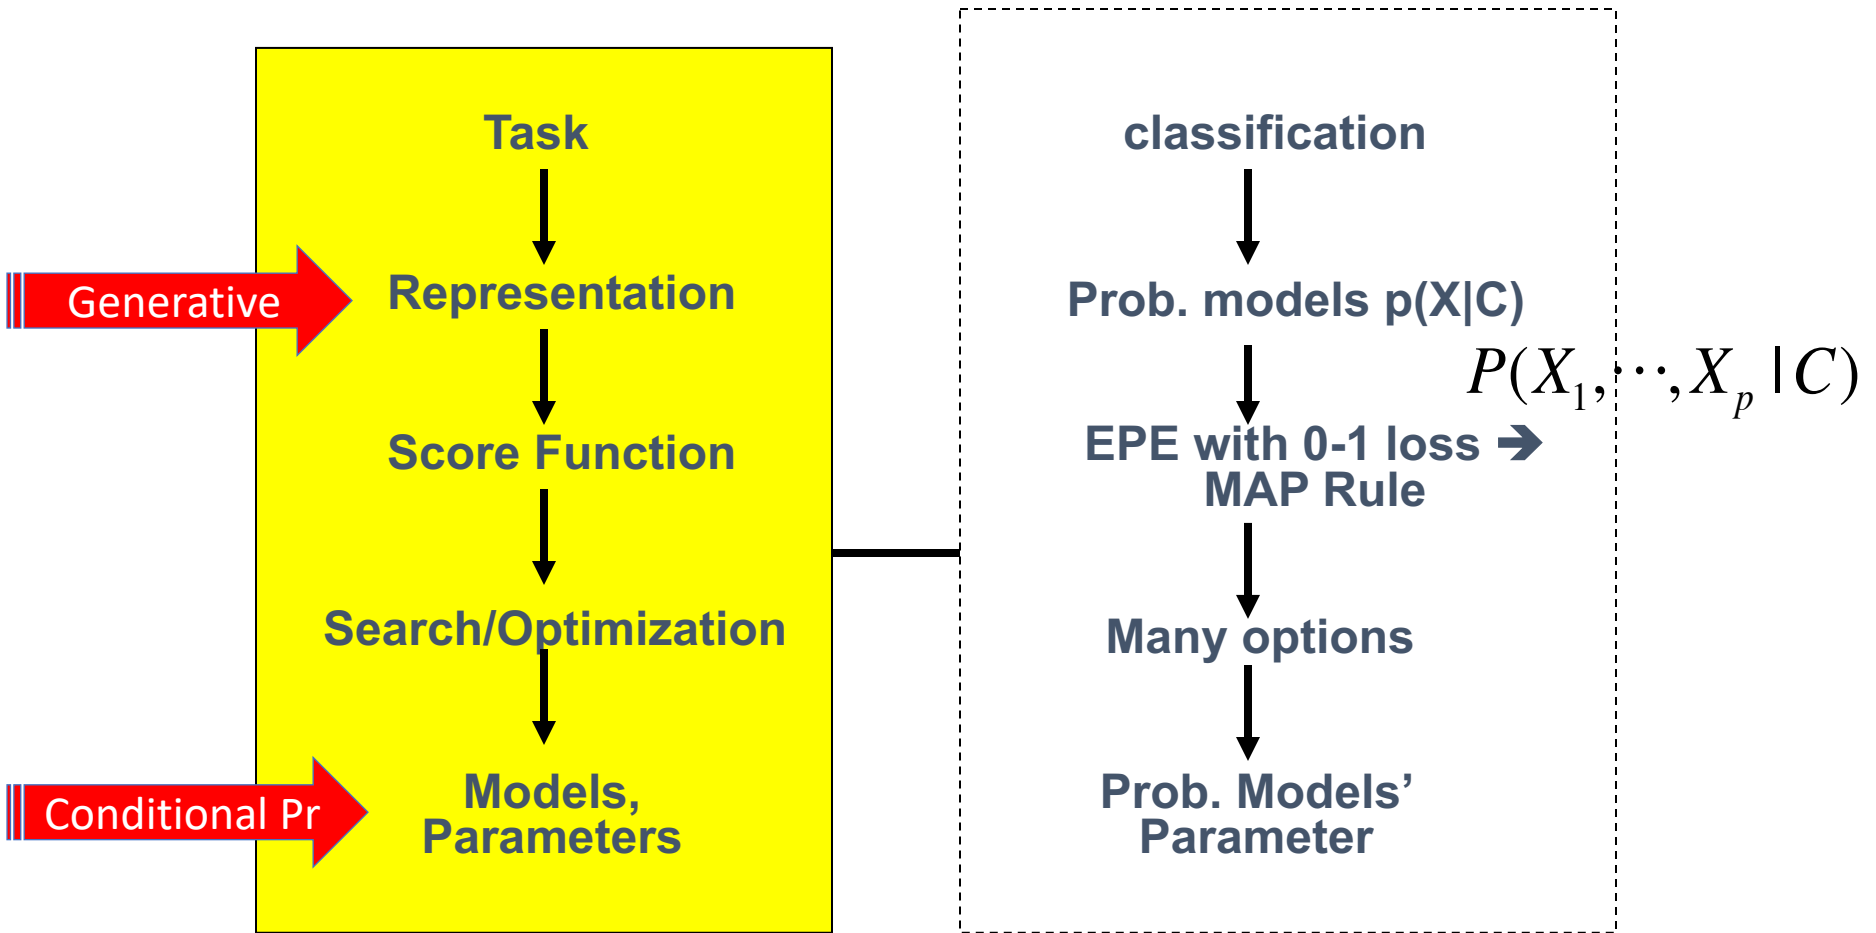
[MAP rule]

$$\operatorname{argmax}_{C_i} P(C_i | X) = \operatorname{argmax}_{C_i} P(X, C_i) = \operatorname{argmax}_{C_i} P(X | C_i) P(C_i)$$



$$\mathbf{X} = (x_1, x_2, \dots, x_p)$$

Generative Bayes Classifier



$$\operatorname{argmax}_k P(C = k | X) = \operatorname{argmax}_k P(X, C) = \operatorname{argmax}_k P(X | C)P(C)$$

X_1	X_2	X_3	C

or
n

$$f : X \longrightarrow C$$

Output as Discrete Class Label
 C_1, C_2, \dots, C_L

Discriminative

$$\operatorname{argmax}_{c \in C} P(c | \mathbf{X}) \quad C = \{c_1, \dots, c_L\}$$

Generative

$$\operatorname{argmax}_{c \in C} P(c | X) = \operatorname{argmax}_{c \in C} P(X, c) = \operatorname{argmax}_{c \in C} P(X | c) P(c)$$

this lecture!

An Example

- Example: Play Tennis

PlayTennis: training examples

Day	Outlook	Temperature	Humidity	Wind	PlayTennis
D1	Sunny ✓	Hot	High	Weak	No
D2	Sunny	Hot	High	Strong	No
D3	Overcast	Hot	High	Weak	Yes
D4	Rain ✓	Mild	High	Weak	Yes
D5	Rain ✓	Cool	Normal	Weak	Yes
D6	Rain	Cool	Normal	Strong	No
D7	Overcast	Cool	Normal	Strong	Yes
D8	Sunny	Mild	High	Weak	No
D9	Sunny	Cool	Normal	Weak	Yes
D10	Rain	Mild	Normal	Weak	Yes
D11	Sunny	Mild	Normal	Strong	Yes
D12	Overcast	Mild	High	Strong	Yes
D13	Overcast	Hot	Normal	Weak	Yes
D14	Rain	Mild	High	Strong	No

X_1 X_2 X_3 X_4 C

X_1	X_2	X_3	C

\mathcal{X}_1	\mathcal{X}_2	\mathcal{X}_3	\mathcal{X}_4	C
S	H	H	W	$C_1 = \text{Yes}$
O	M	N	S	
R	C			$C_2 = \text{No}$
3	* 3	* 2	* 2	* 2 \Rightarrow 72 parameters

$\Rightarrow P(\mathcal{X}_1 \mathcal{X}_2 \mathcal{X}_3 \mathcal{X}_4 | C), P(C)$ training

$$\textcircled{1} \begin{cases} P(C = \text{Yes}) = \frac{N(\text{Yes})}{N(\text{train})} = \frac{9}{14} \\ P(C = \text{No}) = 1 - P(C = \text{Yes}) = \frac{5}{14} \end{cases}$$

$$\textcircled{2} \begin{aligned} P(\mathcal{X}_1 = S, \mathcal{X}_2 = H, \mathcal{X}_3 = H, \mathcal{X}_4 = W | C = \text{No}) &= \frac{N(SHHW\text{No})}{N(\text{No})} = \frac{1}{5} \\ P(\mathcal{X}_1 = S, \mathcal{X}_2 = H, \mathcal{X}_3 = H, \mathcal{X}_4 = W | C = \text{Yes}) &= \frac{0}{9} = 0 \end{aligned}$$

Learning Phase:

- Example: Play Tennis

X_1	X_2	X_3	C

$k_2 = 3$

X_1 ↓ X_2 ↑ X_3 ↑ X_4 ↑ C

PlayTennis: training examples

Day	Outlook	Temperature	Humidity	Wind	PlayTennis
D1	Sunny	Hot	High	Weak	No
D2	Sunny	Hot	High	Strong	No
D3	Overcast	Hot	High	Weak	Yes
D4	Rain	Mild	High	Weak	Yes
D5	Rain	Cool	Normal	Weak	Yes
D6	Rain	Cool	Normal	Strong	No
D7	Overcast	Cool	Normal	Strong	Yes
D8	Sunny	Mild	High	Weak	No
D9	Sunny	Cool	Normal	Weak	Yes
D10	Rain	Mild	Normal	Weak	Yes
D11	Sunny	Mild	Normal	Strong	Yes
D12	Overcast	Mild	High	Strong	Yes
D13	Overcast	Hot	Normal	Weak	Yes
D14	Rain	Mild	High	Strong	No

X_2 : {Hot, Mild, Cool}

$X_3 = \{High, Normal\}$

$k_3 = 2$

$X_4 = (W, S)$
 $k_4 =$

$C: \{Yes, No\}$
 $(L=2)$

$X_1 = \{sunny, overc, rain\}$
 $(k=3)$

- **Learning:** maximum likelihood estimates
 - simply use the frequencies in the data

e.g. $p(\text{overcast, hot, high, weak} \mid \text{Yes}) = \frac{1}{9}$

PlayTennis: training examples

Day	Outlook	Temperature	Humidity	Wind	PlayTennis
D1	Sunny	Hot	High	Weak	No
D2	Sunny	Hot	High	Strong	No
D3	Overcast	Hot	High	Weak	Yes
D4	Rain	Mild	High	Weak	Yes
D5	Rain	Cool	Normal	Weak	Yes
D6	Rain	Cool	Normal	Strong	No
D7	Overcast	Cool	Normal	Strong	Yes
D8	Sunny	Mild	High	Weak	No
D9	Sunny	Cool	Normal	Weak	Yes
D10	Rain	Mild	Normal	Weak	Yes
D11	Sunny	Mild	Normal	Strong	Yes
D12	Overcast	Mild	High	Strong	Yes
D13	Overcast	Hot	Normal	Weak	Yes
D14	Rain	Mild	High	Strong	No

Directly estimate from data

e.g.

$p(\text{overcast, hot, high, weak} \mid \text{No})$

$= \frac{0}{9}$

Check L13-MLE Lecture for Why

$P(X_1, X_2, X_3, X_4 | \text{Yes})$
 $P(X_1, X_2, X_3, X_4 | \text{No})$

36 x 2

X_1	X_2	X_3	C

3 *PlayTennis: training examples*

Day	Outlook	Temperature	Humidity	Wind	PlayTennis
D1	Sunny	Hot	High	Weak	No
D2	Sunny	Hot	High	Strong	No
D3	Overcast	Hot	High	Weak	Yes
D4	Rain	Mild	High	Weak	Yes
D5	Rain	Cool	Normal	Weak	Yes
D6	Rain	Cool	Normal	Strong	No
D7	Overcast	Cool	Normal	Strong	Yes
D8	Sunny	Mild	High	Weak	No
D9	Sunny	Cool	Normal	Weak	Yes
D10	Rain	Mild	Normal	Weak	Yes
D11	Sunny	Mild	Normal	Strong	Yes
D12	Overcast	Mild	High	Strong	Yes
D13	Overcast	Hot	Normal	Weak	Yes
D14	Rain	Mild	High	Strong	No

$P(C=C_i)$

$\left\{ \begin{array}{l} \text{Yes} \\ \text{No} \end{array} \right\}$

$P(C=Yes)$
 $= 9/14$

$P(C=No)$
 $= 5/14$

3 x 3 x 2 x 2 = 36

Generative Bayes Classifier:

e.g English Dictionary $\mathcal{X}_1, \dots, \mathcal{X}_p$

$p \sim 30k$ } $2^p \times L$

$\Rightarrow C$ (up, down) L

$P(C_1), P(C_2), \dots, P(C_L)$

$p(C_1)$

$p(C_2)$

$P(\text{Play}=\text{Yes}) = 9/14$ $P(\text{Play}=\text{No}) = 5/14$

$P(X_1, X_2, \dots, X_p | C_1), P(X_1, X_2, \dots, X_p | C_2)$

a look up table of cond. prob

Outlook (3 values)	Temperature (3 values)	Humidity (2 values)	Wind (2 values)	Play=Yes	Play=No
sunny	hot	high	weak	0/9	1/5
sunny	hot	high	strong	.../9	.../5
sunny	hot	normal	weak	.../9	.../5
sunny	hot	normal	strong	.../9	.../5
....
....
....
....

$3*3*2*2$ [conjunctions of attributes] * 2 [two classes]=72 parameters

Generative Bayes Classifier:

- **Testing Phase**

- – Given an unknown instance $\mathbf{X}'_{ts} = (a'_1, \dots, a'_p)$
- **Look up tables** to assign the label c^* to \mathbf{X}'_{ts} if

$$\hat{P}(a'_1, \dots, a'_p | c^*) \hat{P}(c^*) > \hat{P}(a'_1, \dots, a'_p | c) \hat{P}(c),$$
$$c \neq c^*, c = c_1, \dots, c_L$$

Last
Page

- Given a new instance,

$x' = (\text{Outlook}=\text{Sunny}, \text{Temperature}=\text{Cool}, \text{Humidity}=\text{High}, \text{Wind}=\text{Strong})$

$$\left\{ \begin{array}{l} p(x' | \text{Yes}) p(c = \text{Yes}) \\ p(x' | \text{No}) p(c = \text{No}) \end{array} \right\} \Rightarrow \arg \max_c \Rightarrow \text{predicted } c^*$$

$$P(C = \text{Yes} \mid X_1, X_2, X_3, X_4)$$

$$P(C = \text{No} \mid X_1, X_2, X_3, X_4)$$

$$\rightarrow P(C_1 = \text{Yes}) = 9/14$$

$$P(C_2 = \text{No}) = 5/14$$

$$\rightarrow P(X_1, X_2, X_3, X_4 \mid C_i)$$

$3 \times 3 \times 2 \times 2 \times 2 \Rightarrow 72$ parameters from train

$$\rightarrow \operatorname{argmax}_{i=1,2} P(\bar{X}_{ts} \mid C_i) P(C_i) \quad \text{Generative BC}$$

Today Recap: Generative Bayes Classifiers

- ✓ Bayes Classifier
 - Generative Bayes Classifier
- ✓ Naïve Bayes Classifier
- ✓ Gaussian Bayes Classifiers
 - Gaussian distribution
 - Naïve Gaussian BC
 - Not-naïve Gaussian BC → LDA, QDA

Naïve Bayes Classifier

- Bayes classification

$$\operatorname{argmax}_{c_j \in \mathcal{C}} P(x_1, x_2, \dots, x_p | c_j) P(c_j)$$

Handwritten red notes:
A red underline is drawn under the expression x_1, x_2, \dots, x_p . To the right, the text $2 \times 2 \times \dots \times 2 = 2^p \times L$ is written in red, where L is written above the second 2 in the product.

Difficulty: learning the joint probability

- Naïve Bayes classification
 - Assumption that **all input attributes are conditionally independent!**
(when given c_i)

Naïve Bayes Classifier

- Bayes classification

$$\operatorname{argmax}_{c_j \in \mathcal{C}} \underbrace{P(x_1, x_2, \dots, x_p | c_j)}_{= P(x_1 | c_j) P(x_2 | c_j) \dots P(x_p | c_j)} P(c_j)$$

Difficulty: learning the joint probability

- Naïve Bayes classification

- Assumption that **all input attributes are conditionally independent!**
given C variable

Naïve Bayes Classifier

- Naïve Bayes classification
 - Assumption that **all input attributes are conditionally independent!**

$$P(X_1, X_2, \dots, X_p | C) \stackrel{\text{[assumption]}}{=} P(X_1 | C) P(X_2 | C) \cdots P(X_p | C)$$

$$\begin{aligned} P(C_j | x_1, \dots, x_p) &\propto P(x_1, x_2, \dots, x_p | C_j) P(C_j) \\ &= \underbrace{P(x_1 | C_j)} \underbrace{P(x_2 | C_j)} \cdots \underbrace{P(x_p | C_j)} \underbrace{P(C_j)} \end{aligned}$$

3 x L		$P(\Delta_1 c)$
$\Delta_1 = \text{Sunm}$	C=Yes	$\frac{N(S, \text{Yes})}{N(\text{Yes})} = \frac{2}{9}$
$\Delta_1 = O$	C=Yes	
$\Delta_1 = R$	C=Yes	
$\Delta_1 = S$	C=No	$\frac{N(S, N)}{N(N)} = \frac{3}{5}$
O	C=No	
R	C=No	

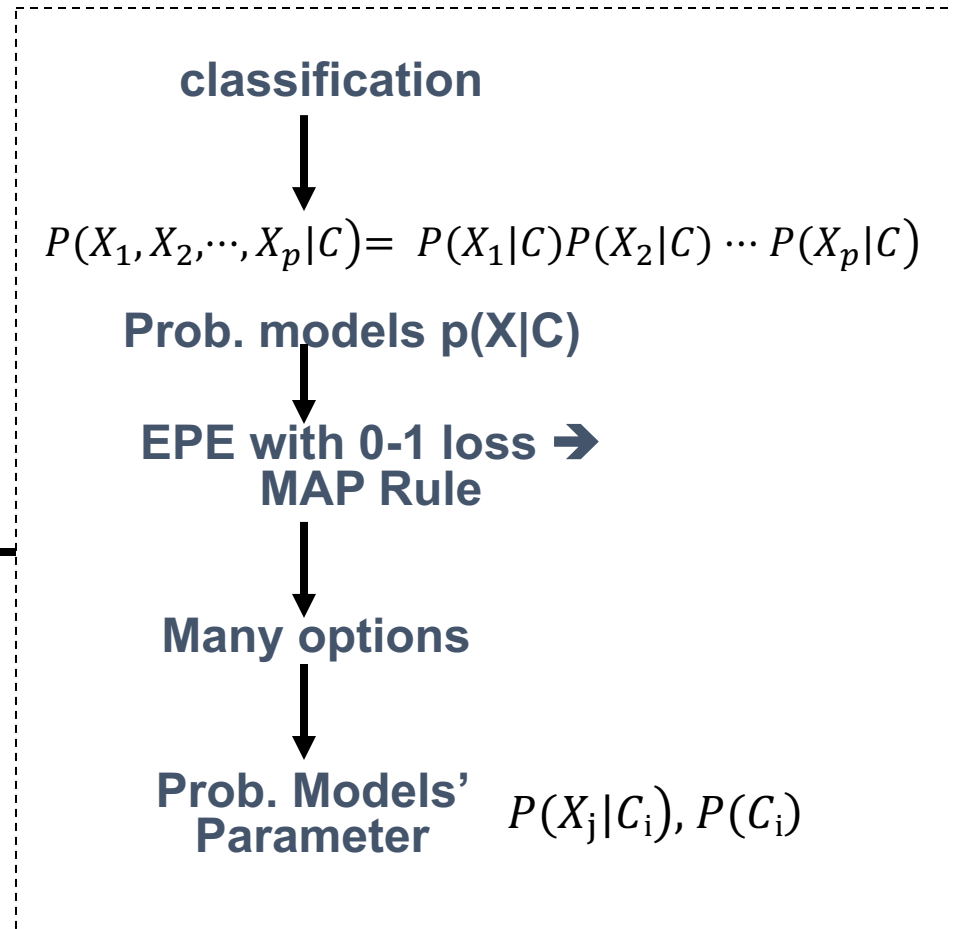
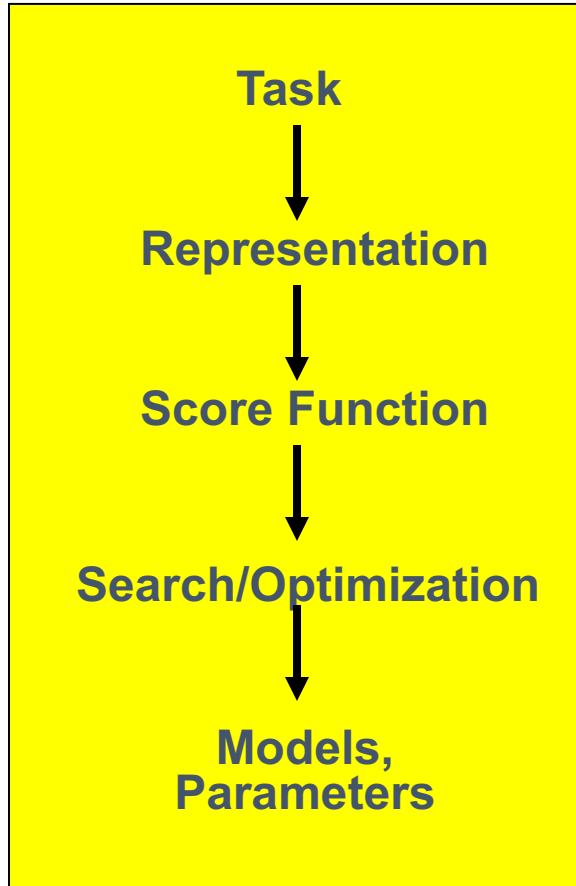
- $\Delta_1 : 3 \times L$
- $\Delta_2 : 3 \times L$
- $\Delta_3 : 2 \times L$
- $\Delta_4 : 2 \times L$

L=2	$P(C=Ye)$	$P(C=N)$
	$\frac{9}{14}$	$\frac{5}{14}$

No Naive
 $72 + 2 = 74$ param
 2^L

$(3 + 3 + 2 + 2) \times L = 20$
 $\Rightarrow \Delta_1 + \Delta_2 + \Delta_3 \Rightarrow 2^L$
 $\underbrace{\hspace{10em}}_{p \sim 30k}$

Naive Bayes Classifier



Estimate $P(X_j = x_{jk} | C = c_i)$ with examples in training;

$P(X_2|C_1), P(X_2|C_2)$

Outlook	Play=Yes	Play=No
<i>Sunny</i>		
<i>Overcast</i>		
<i>Rain</i>		

Temperature	Play=Yes	Play=No
<i>Hot</i>		
<i>Mild</i>		
<i>Cool</i>		

$P(X_4|C_1), P(X_4|C_2)$

Humidity	Play=Yes	Play=No
<i>High</i>		
<i>Normal</i>		

Wind	Play=Yes	Play=No
<i>Strong</i>		
<i>Weak</i>		

$P(\text{Play=Yes}) = ??$

$P(\text{Play=No}) = ??$

$P(C_1), P(C_2), \dots, P(C_L)$

Naïve Bayes Classifier

- Naïve Bayes classification

- Assumption that **all input attributes are conditionally independent!**

$$P(X_1, X_2, \dots, X_p | C) = P(X_1 | C)P(X_2 | C) \cdots P(X_p | C)$$

- MAP classification rule: for a sample $\mathbf{x} = (x_1, x_2, \dots, x_p)$

$$[P(x_1 | c^*) \cdots P(x_p | c^*)]P(c^*) > [P(x_1 | c) \cdots P(x_p | c)]P(c),$$

$$c \neq c^*, c = c_1, \dots, c_L$$

Naïve Bayes Classifier

- Naïve Bayes classification

- Assumption that **all input attributes are conditionally independent!**

$$P(X_1, X_2, \dots, X_p | C) = P(X_1 | C)P(X_2 | C) \cdots P(X_p | C)$$

- MAP classification rule: for a sample $\mathbf{x} = (x_1, x_2, \dots, x_p)$

$$\underbrace{[P(x_1 | c^*) \cdots P(x_p | c^*)]P(c^*)}_{>} > [P(x_1 | c) \cdots P(x_p | c)]P(c),$$

$$c \neq c^*, c = c_1, \dots, c_L$$

$$\Rightarrow \operatorname{argmax}_{i=1, \dots, L} P(c_i)P(x_1|c_i)P(x_2|c_i) \cdots P(x_p|c_i)$$

Naïve Bayes Classifier

- Naïve Bayes classification

- Assumption that **all input attributes are conditionally independent!**

$$P(X_1, X_2, \dots, X_p | C) = P(X_1 | C)P(X_2 | C) \cdots P(X_p | C)$$

- MAP classification rule: for a sample $\mathbf{x} = (x_1, x_2, \dots, x_p)$

$$[P(x_1 | c^*) \cdots P(x_p | c^*)]P(c^*) > [P(x_1 | c) \cdots P(x_p | c)]P(c),$$

$$c \neq c^*, c = c_1, \dots, c_L$$

$\{j=1,2,\dots,p\} \{i=1,2,\dots,L\}$
 $P(X_j | c_i)$

Bernoulli (P)
Binomial (K,P)
Multinomial
Gaussian

Naïve Bayes Classifier (for discrete input attributes) – training/ Learning phase

- Learning Phase: Given a training set S ,

For each target value of c_i ($c_i = c_1, \dots, c_L$)

$\hat{P}(C = c_i) \leftarrow$ estimate $P(C = c_i)$ with examples in S ;

Naïve Bayes Classifier (for discrete input attributes) – training/ Learning phase

- Learning Phase: Given a training set \mathbf{S} ,

For each target value of c_i ($c_i = c_1, \dots, c_L$)

$\hat{P}(C = c_i) \leftarrow$ estimate $P(C = c_i)$ with examples in \mathbf{S} ;

For every attribute value x_{jk} of each attribute X_j ($j = 1, \dots, p$; $k = 1, \dots, K_j$)

$\hat{P}(X_j = x_{jk} | C = c_i) \leftarrow$ estimate $P(X_j = x_{jk} | C = c_i)$ with examples in \mathbf{S} ;

Output: conditional probability tables; for $X_j: K_j \times L$ elements

Naïve Bayes Classifier (for discrete input attributes) - training

- **Learning Phase:** Given a training set S ,

For each target value of c_i ($c_i = c_1, \dots, c_L$)

$\hat{P}(C = c_i) \leftarrow$ estimate $P(C = c_i)$ with examples in S ;

For every attribute value x_{jk} of each attribute X_j ($j = 1, \dots, p$; $k = 1, \dots, K_j$)

$\hat{P}(X_j = x_{jk} | C = c_i) \leftarrow$ estimate $P(X_j = x_{jk} | C = c_i)$ with examples in S ;

$\{k_1, k_2, \dots, k_p\}$
 $\{X_1, X_2, \dots, X_p\}$

$k_1 \times L +$
 $k_2 \times L +$
 $k_p \times L$

Naïve Bayes (for discrete input attributes) - testing

- **Test Phase:** Given an unknown instance


Look up tables to assign the label c^* to \mathbf{X}' if $\mathbf{X}' = (a'_1, \dots, a'_p)$

$$[\hat{P}(a'_1 | c^*) \cdots \hat{P}(a'_p | c^*)] \hat{P}(c^*) > \underbrace{[\hat{P}(a'_1 | c) \cdots \hat{P}(a'_p | c)] \hat{P}(c)},$$


$$c \neq c^*, c = c_1, \dots, c_L$$

$$\begin{aligned} & P(\mathbf{X}' | c_i) P(c_i) \\ &= P(a'_1 | c_i) P(a'_2 | c_i) \cdots P(a'_p | c_i) P(c_i) \\ & \quad i=1, 2, \dots, L \end{aligned}$$

Learning (training) the NBC Model



Learning (training) the NBC Model




- maximum likelihood estimates:
 - simply use the frequencies in the data

$$\hat{P}(c_j) = \frac{N(C = c_j)}{N}$$

$$\hat{P}(x_i | c_j) = \frac{N(X_i = x_i, C = c_j)}{N(C = c_j)}$$


Learning (training) the NBC Model



- maximum likelihood estimates
 - simply use the frequencies in the data

$$\hat{P}(c_j) = \frac{N(C = c_j)}{N}$$

$$\hat{P}(x_i | c_j) = \frac{N(X_i = x_i, C = c_j)}{N(C = c_j)}$$



PlayTennis: training examples

Day	Outlook	Temperature	Humidity	Wind	PlayTennis
D1	Sunny	Hot	High	Weak	No
D2	Sunny	Hot	High	Strong	No
D3	Overcast	Hot	High	Weak	Yes
D4	{ Rain }	Mild	High	Weak	Yes
D5	{ Rain }	Cool	Normal	Weak	Yes
D6	Rain	Cool	Normal	Strong	No
D7	Overcast	Cool	Normal	Strong	Yes
D8	Sunny	Mild	High	Weak	No
D9	Sunny	Cool	Normal	Weak	Yes
D10	{ Rain }	Mild	Normal	Weak	Yes
D11	Sunny	Mild	Normal	Strong	Yes
D12	Overcast	Mild	High	Strong	Yes
D13	Overcast	Hot	Normal	Weak	Yes
D14	Rain	Mild	High	Strong	No

$$P(X_1 = \text{Rain} \mid C = \text{Yes}) = \frac{3}{9}$$

$$P(X_1 = \text{Rain} \mid C = \text{No}) = \frac{2}{5}$$

Counting



• Learning Phase

Estimate $P(X_j = x_{jk} | C = c_i)$ with examples in training;

$P(X_2|C_1), P(X_2|C_2)$

Outlook	Play=Yes	Play=No
Sunny	2/9	3/5
Overcast	4/9	0/5
Rain	3/9	2/5

Temperature	Play=Yes	Play=No
Hot	2/9	2/5
Mild	4/9	2/5
Cool	3/9	1/5

Humidity	Play=Yes	Play=No
High	3/9	4/5
Normal	6/9	1/5

$P(X_4|C_1), P(X_4|C_2)$

Wind	Play=Yes	Play=No
Strong	3/9	3/5
Weak	6/9	2/5

3+3+2+2 [naïve assumption] * 2 [two classes]= 20 parameters

$P(\text{Play=Yes}) = 9/14$ $P(\text{Play=No}) = 5/14$

$P(C_1), P(C_2), \dots, P(C_L)$

Counting

Learning Phase

$P(x_1, x_2, x_3, x_4 | C_i)$

Estimate $P(X_j = x_{jk} | C = c_i)$ with examples in training;

$P(X_2|C_1), P(X_2|C_2)$

X_1
3

Outlook	Play=Yes	Play=No
Sunny	2/9	3/5
Overcast	4/9	0/5
Rain	3/9	2/5

Temperature	Play=Yes	Play=No
Hot	2/9	2/5
Mild	4/9	2/5
Cool	3/9	1/5

2

Humidity	Play=Yes	Play=No
High	3/9	4/5
Normal	6/9	1/5

$P(X_4|C_1), P(X_4|C_2)$

Wind	Play=Yes	Play=No
Strong	3/9	3/5
Weak	6/9	2/5

3+3+2+2 [naïve assumption] * 2 [two classes] = 20 parameters

$P(\text{Play=Yes}) = 9/14$

$P(\text{Play=No}) = 5/14$

$P(C_1), P(C_2), \dots, P(C_L)$

$P(C_i)$

Testing the NBC Model

look up

$$[\hat{P}(a'_1 | c^*) \cdots \hat{P}(a'_p | c^*)] \hat{P}(c^*) > [\hat{P}(a'_1 | c) \cdots \hat{P}(a'_p | c)] \hat{P}(c)$$

- Test Phase

- Given a new instance,

$x' = (\text{Outlook} = \text{Sunny}, \text{Temperature} = \text{Cool}, \text{Humidity} = \text{High}, \text{Wind} = \text{Strong})$

Using the NBC Model

$$[\hat{P}(a'_1 | c^*) \cdots \hat{P}(a'_p | c^*)] \hat{P}(c^*) > [\hat{P}(a'_1 | c) \cdots \hat{P}(a'_p | c)] \hat{P}(c)$$

- Test Phase

- Given a new instance,

$x' = (\text{Outlook} = \text{Sunny}, \text{Temperature} = \text{Cool}, \text{Humidity} = \text{High}, \text{Wind} = \text{Strong})$

$$\begin{aligned} \rightarrow & P(c_1) P(\text{Sunny} | c_1) P(\text{Cool} | c_1) P(\text{High} | c_1) P(\text{Strong} | c_1) \\ &= \frac{9}{14} \times \frac{2}{9} \cdots = \end{aligned}$$

$$\begin{aligned} \rightarrow & P(c_2) P(\text{Su} | c_2) P(\text{Co} | c_2) P(\text{hi} | c_2) P(\text{St} | c_2) \\ &= \frac{5}{14} \times \frac{3}{5} \times \cdots = \end{aligned}$$

Testing the NBC Model

$$[\hat{P}(a'_1 | c^*) \cdots \hat{P}(a'_p | c^*)] \hat{P}(c^*) > [\hat{P}(a'_1 | c) \cdots \hat{P}(a'_p | c)] \hat{P}(c)$$

- Test Phase

- Given a new instance,

$x' = (\text{Outlook}=\textit{Sunny}, \text{Temperature}=\textit{Cool}, \text{Humidity}=\textit{High}, \text{Wind}=\textit{Strong})$

- Look up in conditional-prob tables

$$P(\text{Outlook}=\textit{Sunny} | \text{Play}=\textit{Yes}) = 2/9$$

$$P(\text{Temperature}=\textit{Cool} | \text{Play}=\textit{Yes}) = 3/9$$

$$P(\text{Humidity}=\textit{High} | \text{Play}=\textit{Yes}) = 3/9$$

$$P(\text{Wind}=\textit{Strong} | \text{Play}=\textit{Yes}) = 3/9$$

$$P(\text{Play}=\textit{Yes}) = 9/14$$

$$P(\text{Outlook}=\textit{Sunny} | \text{Play}=\textit{No}) = 3/5$$

$$P(\text{Temperature}=\textit{Cool} | \text{Play}=\textit{No}) = 1/5$$

$$P(\text{Humidity}=\textit{High} | \text{Play}=\textit{No}) = 4/5$$

$$P(\text{Wind}=\textit{Strong} | \text{Play}=\textit{No}) = 3/5$$

$$P(\text{Play}=\textit{No}) = 5/14$$

- Test Phase $[\hat{P}(a'_1 | c^*) \cdots \hat{P}(a'_p | c^*)] \hat{P}(c^*) > [\hat{P}(a'_1 | c) \cdots \hat{P}(a'_p | c)] \hat{P}(c)$

- Given a new instance,

$\mathbf{x}' = (\text{Outlook}=\text{Sunny}, \text{Temperature}=\text{Cool}, \text{Humidity}=\text{High}, \text{Wind}=\text{Strong})$

- Look up in conditional-prob tables

$P(\text{Outlook}=\text{Sunny} | \text{Play}=\text{Yes}) = 2/9$

$P(\text{Temperature}=\text{Cool} | \text{Play}=\text{Yes}) = 3/9$

$P(\text{Humidity}=\text{High} | \text{Play}=\text{Yes}) = 3/9$

$P(\text{Wind}=\text{Strong} | \text{Play}=\text{Yes}) = 3/9$

$P(\text{Play}=\text{Yes}) = 9/14$

$P(\text{Outlook}=\text{Sunny} | \text{Play}=\text{No}) = 3/5$

$P(\text{Temperature}=\text{Cool} | \text{Play}=\text{No}) = 1/5$

$P(\text{Humidity}=\text{High} | \text{Play}=\text{No}) = 4/5$

$P(\text{Wind}=\text{Strong} | \text{Play}=\text{No}) = 3/5$

$P(\text{Play}=\text{No}) = 5/14$

- MAP rule


$P(\text{Yes} | \mathbf{x}')$: $[P(\text{Sunny} | \text{Yes})P(\text{Cool} | \text{Yes})P(\text{High} | \text{Yes})P(\text{Strong} | \text{Yes})]P(\text{Play}=\text{Yes}) = 0.0053$


$P(\text{No} | \mathbf{x}')$: $[P(\text{Sunny} | \text{No}) P(\text{Cool} | \text{No})P(\text{High} | \text{No})P(\text{Strong} | \text{No})]P(\text{Play}=\text{No}) = 0.0206$



Given the fact $P(\text{Yes} | \mathbf{x}') < P(\text{No} | \mathbf{x}')$, we label \mathbf{x}' to be “No”.

WHY ? Naïve Bayes Assumption

- 
- $P(c_j)$
 - Can be estimated from the frequency of classes in the training examples.
 - $P(x_1, x_2, \dots, x_p | c_j)$
 - $O(|X_1| \cdot |X_2| \cdot |X_3| \dots |X_p| \cdot |C|)$ parameters
 - Could only be estimated if a very, very large number of training examples was available.

- 
- $P(x_k | c_j)$
 - $O(|X_1| + |X_2| + |X_3| \dots + |X_p| \cdot |C|)$ parameters
 - Assume that the probability of observing the conjunction of attributes is equal to the product of the individual probabilities $P(x_i | c_j)$.

WHY ? Naïve Bayes Assumption

Assuming $|C| = L$
num of unique values

- $P(c_j)$
 - Can be estimated from the frequency of classes in the training examples.

Assuming $|X_i| = 2, i=1, 2, \dots, P$
 $\Rightarrow 2^P \times L$

Exp

- $P(x_1, x_2, \dots, x_p | c_j)$
 - $O(|X_1| \cdot |X_2| \cdot |X_3| \dots |X_p| \cdot |C|)$ parameters
 - Could only be estimated if a very, very large number of training examples was available.

$(2+2+2+\dots+2) \times L$
 $= 2^P \times L$

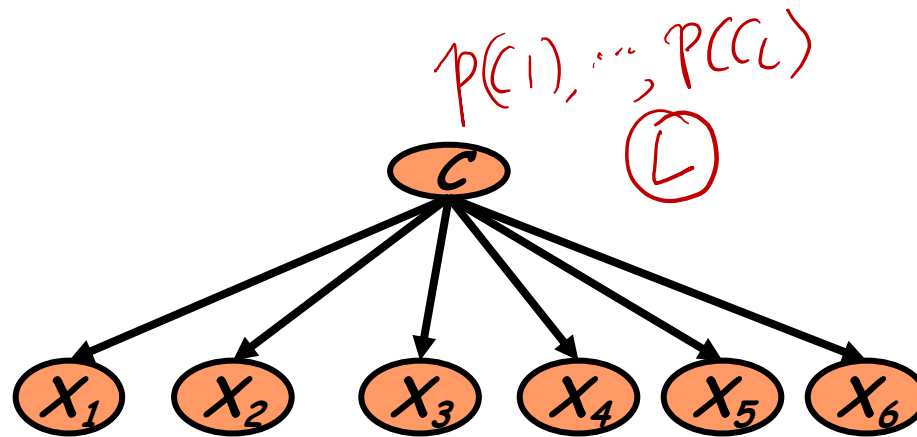
Linear

- $P(x_k | c_j)$
 - $O(|X_1| + |X_2| + |X_3| \dots + |X_p| \cdot |C|)$ parameters
 - Assume that the probability of observing the conjunction of attributes is equal to the product of the individual probabilities $P(x_i | c_j)$.


Not Naïve

Naïve


Challenges during Learning (training) the NBC Model



For instance:



For instance:



- What if we have seen no training cases where patient had no flu and muscle aches?

$$\hat{P}(X_6 = T | C = not_flu) = \frac{N(X_6 = T, C = nf)}{N(C = nf)} = 0$$

muscle-ache-yes/no flu/nf

- Zero probabilities cannot be conditioned away, no matter the other evidence!

$$?? = \arg \max_c \hat{P}(c) \prod_i \hat{P}(x_i | c)$$

$$\delta_f = p(c=f|u) p(x_1|f) p(x_2|f) p(x_3|f) p(x_4|f) p(x_5|f) p(x_6|f)$$

$$\delta_{nf} = p(c=nf) p(x_1/nf) p(x_2/nf) p(x_3/nf) p(x_4/nf) p(x_5/nf) p(x_6/nf)$$

if any term gives 0,

$$\Rightarrow \delta_{nf} = 0$$

no matter other terms' value

Smoothing to Avoid Overfitting

Why necessary ??

$$\hat{P}(x_i | c_j) = \frac{N(X_i = x_i, C = c_j) + 1}{N(C = c_j) + k_i}$$

of values of feature X_i

To make
 $\sum_i (P(x_i | C_j)) = 1$

$$|X_i| = k_i$$

Smoothing to Avoid Overfitting

$$\hat{P}(x_i | c_j) = \frac{N(X_i = x_i, C = c_j) + 1}{N(C = c_j) + k_i}$$

of values of X_i

- Somewhat more subtle version

overall fraction in data
where $X_i = x_{i,k}$

$$\hat{P}(x_{i,k} | c_j) = \frac{N(X_i = x_{i,k}, C = c_j) + mp_{i,k}}{N(C = c_j) + m}$$

$\rightarrow k \in \{1, 2, \dots, k_i\}$

extent of
"smoothing"

Summary:

Generative Bayes Classifier

Task: Classify a new instance X based on a tuple of attribute values $X = \langle X_1, X_2, \dots, X_p \rangle$ into one of the classes

$$c_{MAP} = \operatorname{argmax}_{c_j \in C} P(c_j | x_1, x_2, \dots, x_p)$$

$$= \operatorname{argmax}_{c_j \in C} \frac{P(x_1, x_2, \dots, x_p | c_j) P(c_j)}{P(x_1, x_2, \dots, x_p)}$$

$$= \operatorname{argmax}_{\substack{j=1,2,\dots,L \\ c_j \in C}} \underbrace{P(x_1, x_2, \dots, x_p | c_j)} \underbrace{P(c_j)}$$

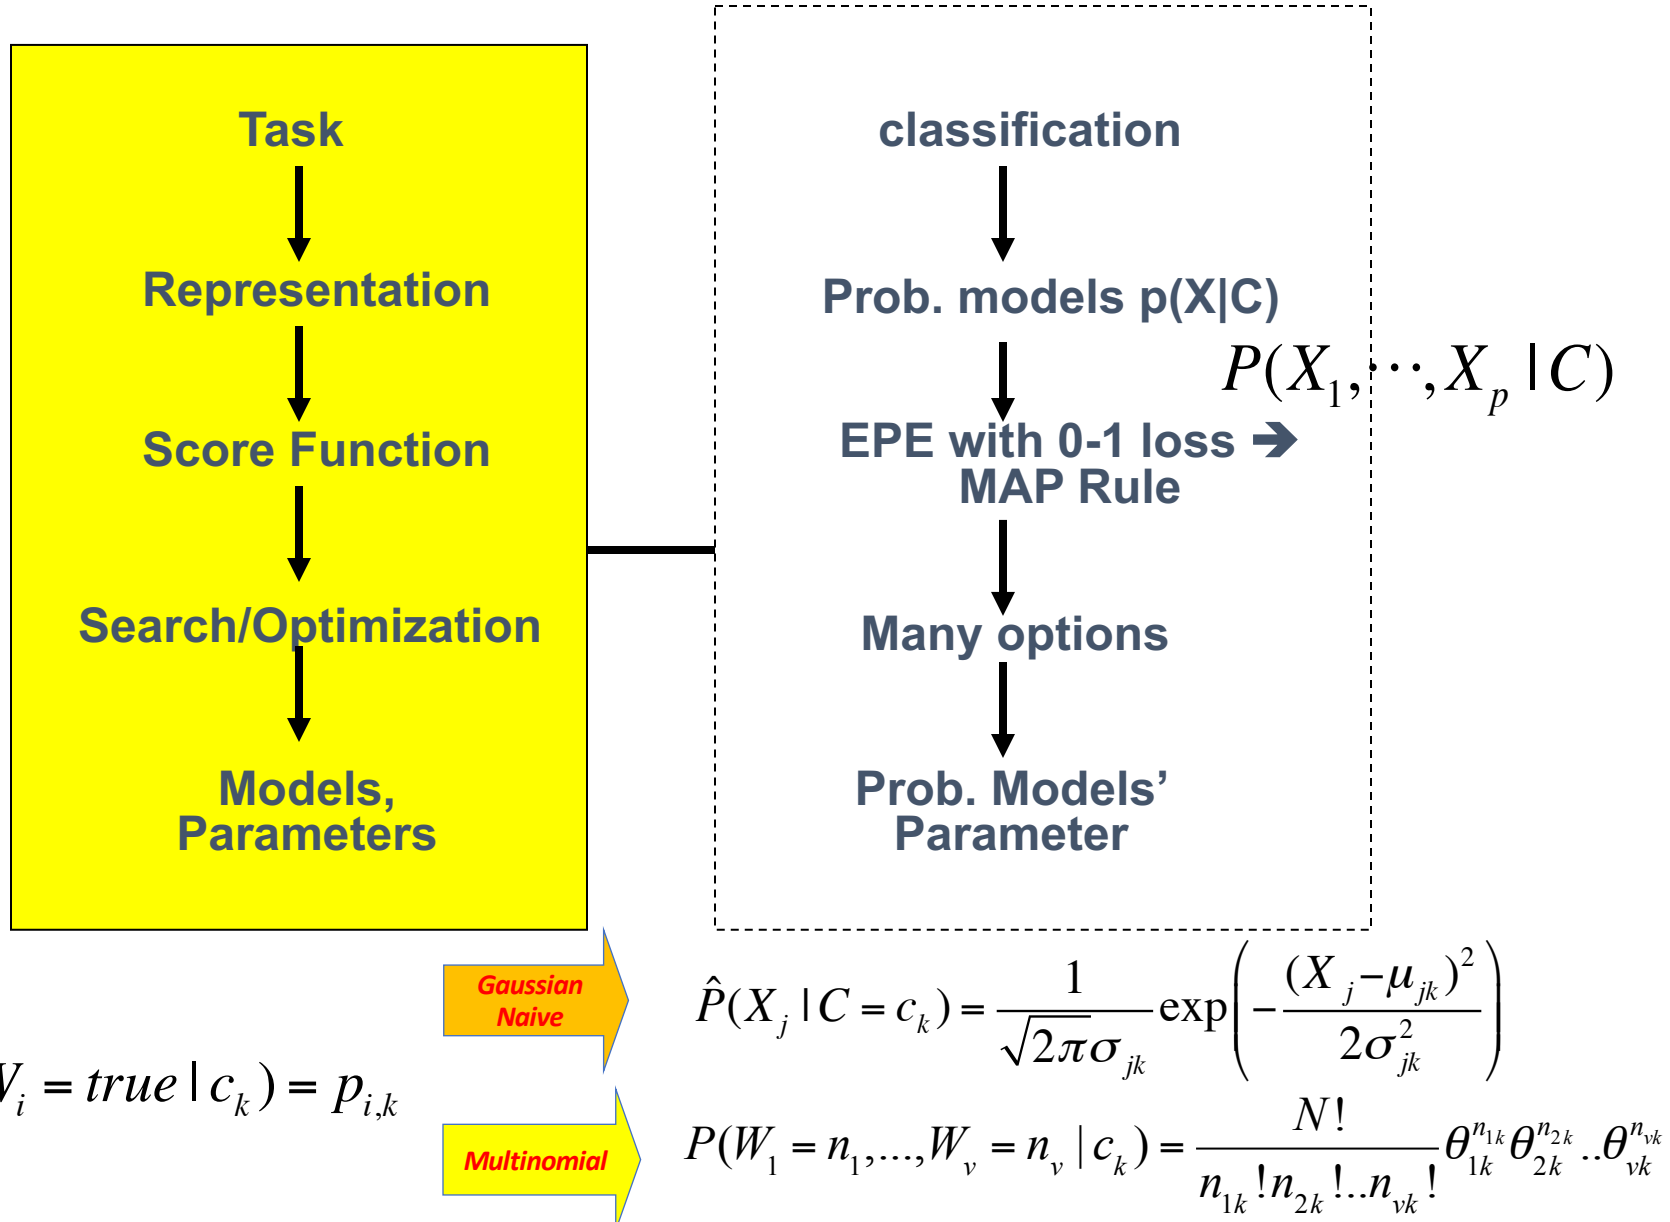
MAP = Maximum A Posteriori

NEXT: More Generative Bayes Classifiers

- ✓ Generative Bayes Classifier
- ✓ Naïve Bayes Classifier
- ✓ Gaussian Bayes Classifiers
 - Gaussian distribution
 - Naïve Gaussian BC
 - Not-naïve Gaussian BC → LDA, QDA
- ✓ Discriminative vs. Generative

$$\operatorname{argmax}_k P(C = k | X) = \operatorname{argmax}_k P(X, C) = \operatorname{argmax}_k P(X | C)P(C)$$

Generative Bayes Classifier



References

- Prof. Andrew Moore's review tutorial
- Prof. Ke Chen NB slides
- Prof. Carlos Guestrin recitation slides