

CS540 Introduction to Artificial Intelligence

Lecture 8

Young Wu

Based on lecture slides by Jerry Zhu, Yingyu Liang, and Charles Dyer

June 20, 2022

Discriminative Model vs Generative Model

Motivation

Generative Models

Motivation

- In probability terms, discriminative models are estimating $\mathbb{P}\{Y|X\}$, the conditional distribution. For example, $a_i \approx \mathbb{P}\{y_i = 1|x_i\}$ and $1 - a_i \approx \mathbb{P}\{y_i = 0|x_i\}$.
- Generative models are estimating $\mathbb{P}\{Y, X\}$, the joint distribution.
- Bayes rule is used to perform classification tasks.

$$\mathbb{P}\{Y|X\} = \frac{\mathbb{P}\{Y, X\}}{\mathbb{P}\{X\}} = \frac{\mathbb{P}\{X|Y\} \mathbb{P}\{Y\}}{\mathbb{P}\{X\}}$$

Joint Distribution

Motivation

- The joint distribution of X_j and $X_{j'}$ provides the probability of $X_j = x_j$ and $X_{j'} = x_{j'}$ occur at the same time.

$$\mathbb{P} \{X_j = x_j, X_{j'} = x_{j'}\}$$

- The marginal distribution of X_j can be found by summing over all possible values of $X_{j'}$.

$$\mathbb{P} \{X_j = x_j\} = \sum_{x \in X_{j'}} \mathbb{P} \{X_j = x_j, X_{j'} = x\}$$

Conditional Distribution

Motivation

- Suppose the joint distribution is given.

$$\mathbb{P}\{X_j = x_j, X_{j'} = x_{j'}\}$$

- The conditional distribution of X_j given $X_{j'} = x_{j'}$ is ratio between the joint distribution and the marginal distribution.

$$\mathbb{P}\{X_j = x_j | X_{j'} = x_{j'}\} = \frac{\mathbb{P}\{X_j = x_j, X_{j'} = x_{j'}\}}{\mathbb{P}\{X_{j'} = x_{j'}\}}$$

Bayes Rule Example 1 Distribution

Quiz

Bayesian Network

Definition

- A Bayesian network is a directed acyclic graph (DAG) and a set of conditional probability distributions.
- Each vertex represents a feature X_j .
- Each edge from X_j to $X_{j'}$ represents that X_j directly influences $X_{j'}$.
- No edge between X_j and $X_{j'}$ implies independence or conditional independence between the two features.

Common Effect

Definition

- For three events A, B, C , the configuration $A \rightarrow B \leftarrow C$ is called common effect.
- In this configuration, A is independent of C , but A is not conditionally independent of C given information about B .
- Once B is observed, A and C are not independent.

Training Bayes Net

Definition

- Training a Bayesian network given the DAG is estimating the conditional probabilities. Let $P(X_j)$ denote the parents of the vertex X_j , and $p(X_j)$ be realizations (possible values) of $P(X_j)$.

$$\mathbb{P}\{x_j | p(X_j)\}, p(X_j) \in P(X_j)$$

- It can be done by maximum likelihood estimation given a training set.

$$\hat{\mathbb{P}}\{x_j | p(X_j)\} = \frac{c_{x_j, p(X_j)}}{c_{p(X_j)}}$$

Bayesian Network Diagram

Quiz

- Story: either Amber (H) or Johnny's dog (D) stepped on a bee, and put something on Johnny's bed (B), and given there is something on Johnny's bed (B), Johnny (J) and Amber (A) can be unhappy.

H	D	B	J	A
0	0	0	1	0
0	0	0	0	1
1	0	0	1	1
0	1	0	1	1
0	0	1	1	0
0	0	1	0	1
1	0	1	1	0
0	1	1	1	0

Bayesian Network Diagram CPT Count

Quiz

Bayes Net Training Example, Training Quiz

- Given a network and the training data.
 $H \rightarrow B, D \rightarrow B, B \rightarrow J, B \rightarrow A.$

H	D	B	J	A
0	0	0	1	0
0	0	0	0	1
1	0	0	1	1
0	1	0	1	1
0	0	1	1	0
0	0	1	0	1
1	0	1	1	0
0	1	1	1	0

Bayes Net Training Example, Training 1

Quiz

- Compute $\hat{\mathbb{P}}\{D = 1\}$.

<i>H</i>	<i>D</i>	<i>B</i>	<i>J</i>	<i>A</i>
0	0	0	1	0
0	0	0	0	1
1	0	0	1	1
0	1	0	1	1
0	0	1	1	0
0	0	1	0	1
1	0	1	1	0
0	1	1	1	0

Bayes Net Training Example, Training 2

Quiz

- Compute $\hat{\mathbb{P}}\{J = 1|B = 1\}$.

H	D	B	J	A
0	0	0	1	0
0	0	0	0	1
1	0	0	1	1
0	1	0	1	1
0	0	1	1	0
0	0	1	0	1
1	0	1	1	0
0	1	1	1	0

Bayes Net Training Example, Training 3

Quiz

- What is the conditional probability $\hat{\mathbb{P}}\{J = 1|B = 0\}$?
- A: I don't understand, B: $\frac{1}{4}$, C: $\frac{1}{2}$, D: $\frac{3}{4}$, E: 1

<i>H</i>	<i>D</i>	<i>B</i>	<i>J</i>	<i>A</i>
0	0	0	1	0
0	0	0	0	1
1	0	0	1	1
0	1	0	1	1
0	0	1	1	0
0	0	1	0	1
1	0	1	1	0
0	1	1	1	0

Bayes Net Training Example, Training 4

Quiz

- Compute $\hat{\mathbb{P}}\{B = 1 | H = 0, D = 1\}$.

H	D	B	J	A
0	0	0	1	0
0	0	0	0	1
1	0	0	1	1
0	1	0	1	1
0	0	1	1	0
0	0	1	0	1
1	0	1	1	0
0	1	1	1	0

Bayes Net Training Example, Training 5

Quiz

- What is the conditional probability $\hat{\mathbb{P}}\{B = 1 | H = 0, D = 0\}$?
- A: I don't understand, B: $\frac{1}{4}$, C: $\frac{1}{2}$, D: $\frac{3}{4}$, E: 1

<i>H</i>	<i>D</i>	<i>B</i>	<i>J</i>	<i>A</i>
0	0	0	1	0
0	0	0	0	1
1	0	0	1	1
0	1	0	1	1
0	0	1	1	0
0	0	1	0	1
1	0	1	1	0
0	1	1	1	0

Bayes Net Training Example, Training 5

Quiz

- What is the conditional probability $\hat{\mathbb{P}}\{A = 0 | H = 1, D = 1\}$?
- A : I don't understand, B: 0 , C: $\frac{1}{2}$, D: 1 , E: NA

<i>H</i>	<i>D</i>	<i>B</i>	<i>J</i>	<i>A</i>
0	0	0	1	0
0	0	0	0	1
1	0	0	1	1
0	1	0	1	1
0	0	1	1	0
0	0	1	0	1
1	0	1	1	0
0	1	1	1	0

Laplace Smoothing

Definition

- Recall that the MLE estimation can incorporate Laplace smoothing.

$$\hat{\mathbb{P}}\{x_j | p(X_j)\} = \frac{c_{x_j, p(X_j)} + 1}{c_{p(X_j)} + |X_j|}$$

- Here, $|X_j|$ is the number of possible values (number of categories) of X_j .
- Laplace smoothing is considered regularization for Bayesian networks because it avoids overfitting the training data.

Bayes Net Inference 1

Definition

- Given the conditional probability table, the joint probabilities can be calculated using conditional independence.

$$\begin{aligned}\mathbb{P}\{x_1, x_2, \dots, x_m\} &= \prod_{j=1}^m \mathbb{P}\{x_j | x_1, x_2, \dots, x_{j-1}, x_{j+1}, \dots, x_m\} \\ &= \prod_{j=1}^m \mathbb{P}\{x_j | p(x_j)\}\end{aligned}$$

Bayes Net Inference 2

Definition

- Given the joint probabilities, all other marginal and conditional probabilities can be calculated using their definitions.

$$\mathbb{P}\{x_j | x_{j'}, x_{j''}, \dots\} = \frac{\mathbb{P}\{x_j, x_{j'}, x_{j''}, \dots\}}{\mathbb{P}\{x_{j'}, x_{j''}, \dots\}}$$

$$\mathbb{P}\{x_j, x_{j'}, x_{j''}, \dots\} = \sum_{x_k: k \neq j, j', j'', \dots} \mathbb{P}\{x_1, x_2, \dots, x_m\}$$

$$\mathbb{P}\{x_{j'}, x_{j''}, \dots\} = \sum_{x_k: k \neq j', j'', \dots} \mathbb{P}\{x_1, x_2, \dots, x_m\}$$

Bayes Net Inference Example 1

Quiz

- Assume the network is trained on a larger set with the following CPT. Compute $\hat{\mathbb{P}}\{H = 0, D = 1 | J = 1, A = 0\}$?

$$\hat{\mathbb{P}}\{H = 1\} = 0.001, \hat{\mathbb{P}}\{D = 1\} = 0.001$$

$$\hat{\mathbb{P}}\{B = 1 | H = 1, D = 1\} = 0.95, \hat{\mathbb{P}}\{B = 1 | H = 1, D = 0\} = 0.94$$

$$\hat{\mathbb{P}}\{B = 1 | H = 0, D = 1\} = 0.29, \hat{\mathbb{P}}\{B = 1 | H = 0, D = 0\} = 0.00$$

$$\hat{\mathbb{P}}\{J = 1 | B = 1\} = 0.9, \hat{\mathbb{P}}\{J = 1 | B = 0\} = 0.05$$

$$\hat{\mathbb{P}}\{A = 1 | B = 1\} = 0.7, \hat{\mathbb{P}}\{A = 1 | B = 0\} = 0.01$$

Bayes Net Inference Example 1 Computation 1

Quiz

Bayes Net Inference Example 1 Computation 2

Quiz

Bayes Net Inference Example 2

Quiz

- Compute $\hat{\mathbb{P}}\{D = 1|H = 0\}$?

$$\hat{\mathbb{P}}\{H = 1\} = 0.001, \hat{\mathbb{P}}\{D = 1\} = 0.001$$

$$\hat{\mathbb{P}}\{B = 1|H = 1, D = 1\} = 0.95, \hat{\mathbb{P}}\{B = 1|H = 1, D = 0\} = 0.94$$

$$\hat{\mathbb{P}}\{B = 1|H = 0, D = 1\} = 0.29, \hat{\mathbb{P}}\{B = 1|H = 0, D = 0\} = 0.00$$

- A : 0, B: 0.001, C: 0.0094, D: 0.0095, E: 1

Bayes Net Inference Example 2 Derivation

Quiz

Bayes Net Inference Example 3

Quiz

- Compute $\hat{\mathbb{P}}\{H = 0, D = 1 | B = 1\}$?

$$\hat{\mathbb{P}}\{H = 1\} = 0.001, \hat{\mathbb{P}}\{D = 1\} = 0.001$$

$$\hat{\mathbb{P}}\{B = 1 | H = 1, D = 1\} = 0.95, \hat{\mathbb{P}}\{B = 1 | H = 1, D = 0\} = 0.94$$

$$\hat{\mathbb{P}}\{B = 1 | H = 0, D = 1\} = 0.29, \hat{\mathbb{P}}\{B = 1 | H = 0, D = 0\} = 0.00$$

- A : 0, B: 0.001, C: 0.0094, D: 0.0095, E: 1

Bayes Net Inference Example 3 Derivation

Quiz

Bayes Net Inference Example 4

Quiz

- Compute $\hat{\mathbb{P}}\{B = 1|J = 1, A = 0\}$?

$$\hat{\mathbb{P}}\{J = 1|B = 1\} = 0.9, \hat{\mathbb{P}}\{J = 1|B = 0\} = 0.05$$

$$\hat{\mathbb{P}}\{A = 1|B = 1\} = 0.7, \hat{\mathbb{P}}\{A = 1|B = 0\} = 0.01$$

Given

$$\mathbb{P}\{B = 1\} = 0.001 \cdot 0.001 \cdot 0.95 + 0.001 \cdot 0.999 \cdot (0.94 + 0.29).$$

- A : 0, B: 0.001, C: 0.0094, D: 0.0095, E: 1

Bayes Net Inference Example 4 Derivation

Quiz

Network Structure

Discussion

- Selecting from all possible structures (DAGs) is too difficult.
- Usually, a Bayesian network is learned with a tree structure.
- Choose the tree that maximizes the likelihood of the training data.

Chow Liu Algorithm

Discussion

- Add an edge between features X_j and $X_{j'}$ with edge weight equal to the information gain of X_j given $X_{j'}$ for all pairs j, j' .
- Find the maximum spanning tree given these edges. The spanning tree is used as the structure of the Bayesian network.

Classification Problem

Discussion

- Bayesian networks do not have a clear separation of the label Y and the features X_1, X_2, \dots, X_m .
- The Bayesian network with a tree structure and Y as the root and X_1, X_2, \dots, X_m as the leaves is called the Naive Bayes classifier.
- Bayes rules is used to compute $\mathbb{P}\{Y = y|X = x\}$, and the prediction \hat{y} is y that maximizes the conditional probability.

$$\hat{y}_i = \operatorname{argmax}_y \mathbb{P}\{Y = y|X = x_i\}$$

Naive Bayes Diagram

Discussion

Gaussian Naive Bayes

Discussion

- If the features are not categorical, continuous distributions can be estimated using MLE as the conditional distribution.
- Gaussian Naive Bayes is used if $X_j|Y = y$ is assumed to have the normal distribution.

$$\lim_{\epsilon \rightarrow 0} \frac{1}{\epsilon} \mathbb{P} \{x < X_j \leq x + \epsilon | Y = y\} = \frac{1}{\sqrt{2\pi}\sigma_y^{(j)}} \exp \left(-\frac{(x - \mu_y^{(j)})^2}{2(\sigma_y^{(j)})^2} \right)$$

Gaussian Naive Bayes Training

Discussion

- Training involves estimating $\mu_y^{(j)}$ and $\sigma_y^{(j)}$ since they completely determine the distribution of $X_j | Y = y$.
- The maximum likelihood estimates of $\mu_y^{(j)}$ and $\left(\sigma_y^{(j)}\right)^2$ are the sample mean and variance of the feature j .

$$\hat{\mu}_y^{(j)} = \frac{1}{n_y} \sum_{i=1}^n x_{ij} \mathbb{1}_{\{y_i=y\}}, \quad n_y = \sum_{i=1}^n \mathbb{1}_{\{y_i=y\}}$$

$$\left(\hat{\sigma}_y^{(j)}\right)^2 = \frac{1}{n_y} \sum_{i=1}^n \left(x_{ij} - \hat{\mu}_y^{(j)}\right)^2 \mathbb{1}_{\{y_i=y\}}$$

$$\text{sometimes } \left(\hat{\sigma}_y^{(j)}\right)^2 \approx \frac{1}{n_y - 1} \sum_{i=1}^n \left(x_{ij} - \hat{\mu}_y^{(j)}\right)^2 \mathbb{1}_{\{y_i=y\}}$$

