Learning From Data
Lecture 4: Generative Learning Algorithms

Yang Li **yangli@sz.tsinghua.edu.cn**

October 9, 2021

Introduction

Today's Lecture

Supervised Learning (Part III)

- Discriminative & Generative Models
- Gaussian Discriminant Analysis

Discriminative & Generative Models

Two Learning Approaches

Classify input data x into two classes $y \in \{0, 1\}$

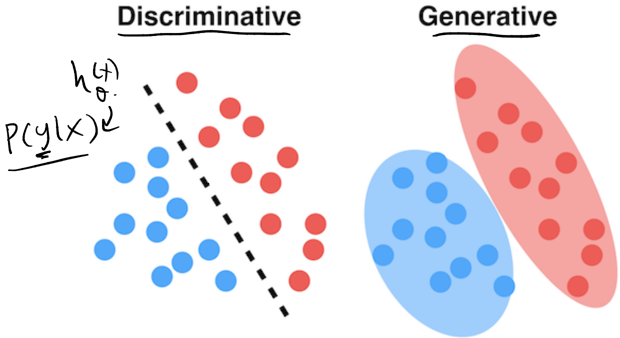

Discriminate between classes of data points

Model the underlying distribution of the data

$P(x, y)$
joint distribution

$$X \xrightarrow{h_\theta(x)} Y$$

Discriminative Learning Algorithms

A class of learning algorithms that try to learn the **conditional probability** $p(y|x)$ directly or learn mappings directly from $\mathcal{X}$ to $\mathcal{Y}$.

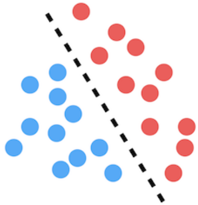

- e.g. linear regression, logistic regression, k-Nearest Neighbors ...

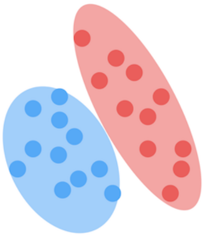

Generative Learning Algorithms

A class of learning algorithms that model the **joint probability** $p(x, y)$ $= p(x|y)p(y)$

- Equivalently, generative algorithms model $p(x|y)$ and $p(y)$
- $p(y)$ is called the **class prior**
- Learned models are transformed to $p(y|x)$ later to classify data using Bayes' rule

Bayes Rule

The posterior distribution on y given x:

$y \in \{0,1\}$.

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

$p(x|y=0)p(0) + p(x|y=1)p(1)$

Bayes Rule

evidence

The posterior distribution on y given x:

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

Make predictions in a generative model:

$$\begin{aligned} \operatorname*{argmax}_y p(y|x) &= \operatorname*{argmax}_y \left(\frac{p(x|y)p(y)}{p(x)} \right) \text{ doesn't depend on } y \\ &= \operatorname*{argmax}_y p(x|y)p(y) \end{aligned}$$

No need to calculate $p(x)$.

Generative Models

Generative classification algorithms:

- Continuous input: Gaussian Discriminant Analysis (GDA)
- Discrete input: Naïve Bayes

Gaussian Discriminant Analysis

Gaussian Discriminant Analysis: Overview

$$\phi \to y \xrightarrow{x|y} x$$

$$\begin{cases} N(\mu_0, \Sigma) & \text{if } y = 0 \\ N(\mu_1, \Sigma) & \text{if } y = 1 \end{cases}$$

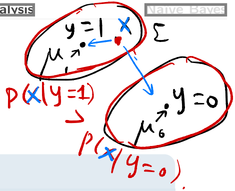

Goal

Binary classification with input in $\mathcal{X} = \mathbb{R}^n$ and label in $\mathcal{Y} = \{0, 1\}$

Main steps

1. Select a *data generating distribution* .

$$y \sim Bernoulli(\phi) \qquad p(y).$$
$$x|y = 0 \sim N(\mu_0, \Sigma), x|y = 1 \sim N(\mu_1, \Sigma)$$

2. Estimate model parameters ϕ, μ_0 ,μ_1 and Σ from training data. $p(x, y)$
3. For any new sample x', predict its label by computing $p(y|x = x'; \phi, \mu_0, \mu_1, \Sigma)$

Multivariate Normal Distribution

Multivariate normal (or multivariate Gaussian) distribution $N(\mu, \Sigma)$

$n \times 1$ $n \times n$

- $\mu \in \mathbb{R}^n$ is the mean vector,
- $\Sigma \in \mathbb{R}^{n \times n}$ is the covariance matrix. Σ is symmetric and SPD.

Density function:

$$p(x; \mu, \Sigma) = \frac{1}{(2\pi)^{n/2} |\Sigma|^{1/2}} e^{\left(-\frac{1}{2}(x-\mu)^T \Sigma^{-1} (x-\mu)\right)}$$

determinant of Σ

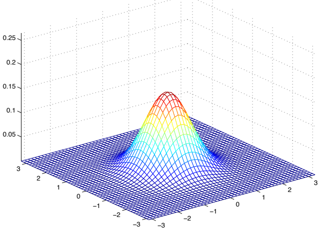

Multivariate Normal Distribution

Let $X \in \mathbb{R}^n$ be a random vector. If $X \sim N(\mu, \Sigma)$,

$$\mathbb{E}[X] = \int_x p(x; \mu, \Sigma) dx = \mu$$

$$\text{Cov}(X) = \mathbb{E}\left[(X - \mathbb{E}[X])(X - \mathbb{E}[X])^T\right] = \Sigma$$

Gaussian Discriminative Analysis

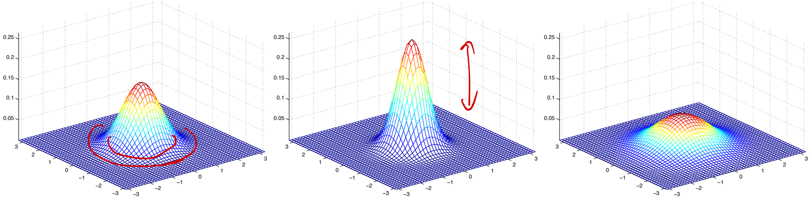

$$\Sigma = \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix} \qquad \Sigma = \begin{bmatrix} 0.6 & 0 \\ 0 & 0.6 \end{bmatrix} \qquad \Sigma = \begin{bmatrix} 2 & 0 \\ 0 & 2 \end{bmatrix}$$

Diagonal entries of Σ controls the "spread" of the distribution

Gaussian Discriminative Analysis

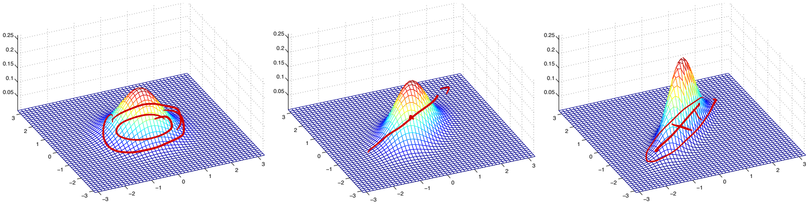

$$\Sigma = \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix} \qquad \Sigma = \begin{bmatrix} 1 & 0.5 \\ 0.5 & 1 \end{bmatrix} \qquad \Sigma = \begin{bmatrix} 1 & 0.8 \\ 0.8 & 1 \end{bmatrix}$$

The distribution is no longer oriented along the axes when off-diagonal entries of Σ are non-zero.

Gaussian Discriminant Analysis (GDA) Model

Given parameters $\phi, \mu_0, \mu_1, \Sigma$,

$$\underline{y \sim \text{Bernoulli}(\phi)}$$

$$\left.\begin{aligned} x|y = 0 &\sim \mathcal{N}(\mu_0, \Sigma) \\ x|y = 1 &\sim \mathcal{N}(\mu_1, \Sigma) \end{aligned}\right\}$$

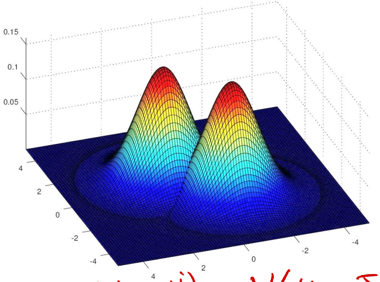

$p(x|y=y^i) \sim \mathcal{N}(\mu_{y_i}, \Sigma)$

Probability density functions:

$$p(y) = \phi^y (1-\phi)^{1-y}$$

$$\left.\begin{aligned} p(x|y=0) &= \frac{1}{(2\pi)^{n/2} |\Sigma|^{1/2}} e^{\left(-\frac{1}{2}(x-\mu_0)^T \Sigma^{-1} (x-\mu_0)\right)} \\ p(x|y=1) &= \frac{1}{(2\pi)^{n/2} |\Sigma|^{1/2}} e^{\left(-\frac{1}{2}(x-\mu_1)^T \Sigma^{-1} (x-\mu_1)\right)} \end{aligned}\right\}$$

Log likelihood of the data:

$p(y|x)$

GDA: $\Sigma = \Sigma_1 = \Sigma_2$

QDA: $\Sigma_1 \neq \Sigma_2$

$$l(\phi, \mu_0, \mu_1, \Sigma) = \log \prod_{i=1}^{m} \underline{p(x^{(i)}, y^{(i)}; \phi, \mu_0, \mu_1, \Sigma)}$$

class prior

$$= \log \prod_{i=1}^{m} \underline{p(x^{(i)}|y^{(i)}; \mu_0, \mu_1, \Sigma)}\, \underline{p(y^{(i)}; \phi)}$$

$$= \sum_{i=1}^{m} \log \underline{P(x^i|y^i;\mu_0,\mu_1,\Sigma)} + \sum_{i=1}^{m} \log P(y^{(i)};\phi)$$

① $P(x|y;\mu_0,\mu_1,\Sigma)$

$$= \underline{\frac{1}{(2\pi)^{\frac{n}{2}}|\Sigma|^{\frac{1}{2}}} e^{-\frac{1}{2}(x-\mu_y)^T\Sigma^{-1}(x-\mu_y)}}$$

② $P(y;\phi) = \underline{\phi^{y}(1-\phi)^{1-y}}$

$$L = \sum_{i=1}^{m} \log \frac{1}{(2\pi)^{\frac{n}{2}}|\Sigma|^{\frac{1}{2}}} + \left(-\frac{1}{2}(x^i-\mu_{y^i})^T\Sigma^{-1}(x^i-\mu_{y^i})\right)$$

$$+ \sum_{i=1}^{m} y^i \log\phi + (1-y^i)\log(1-\phi)$$

Next step:

$\frac{\partial L}{\partial \phi} = 0$, solve for ϕ

$\frac{\partial L}{\partial \mu_0} = 0$, $\frac{\partial L}{\partial \mu_1} = 0$, $\frac{\partial L}{\partial \Sigma} = 0$, solve for μ_0, μ_1, Σ

$$L = \sum_{i=1}^{m} \log \frac{1}{(2\pi)^{\frac{n}{2}}|\Sigma|^{\frac{1}{2}}} + \left(-\frac{1}{2}(x^i-\mu_{y^i})^T\Sigma^{-1}(x^i-\mu_{y^i})\right)$$
$$+ \sum_{i=1}^{m} y^i \log\phi + (1-y^i)\log(1-\phi)$$

Facts

① $\nabla_x(x^TAx)$

$= Ax + A^Tx$

A is symmetric

$= 2Ax.$

$$\frac{\partial L}{\partial \phi} = \sum_{i=1}^{m} y^{(i)} \frac{\partial}{\partial\phi}\log\phi + (1-y^{(i)})\frac{\partial}{\partial\phi}\log(1-\phi)$$

$$= \sum_{i=1}^{m} \frac{y^{(i)}}{\phi} + \frac{(1-y^{(i)})}{1-\phi}(-1)$$

$$= \frac{1}{\phi}\sum_{i=1}^{m} y^{(i)} - \frac{1}{1-\phi}\sum_{i=1}^{m}(1-y^{(i)})$$

$$c := \left[\begin{array}{l}\#\text{ of 1's in}\\ \text{training data}\end{array}\right] = \sum_{i=1}^{m} 1\{y^{(i)}=1\}.$$

$$= \frac{c}{\phi} - \frac{m-c}{1-\phi} = 0$$

$$c(1-\phi) = \phi(m-c)$$

$$c - c\phi - \phi(m-c) = 0.$$

$$c - \phi(c+m-c) = 0.$$

$$\phi = \frac{c}{m} = \frac{1}{m}\sum_{i=1}^{m}(y^{(i)}=1)$$

$$\frac{\partial L}{\partial \mu_0} = \sum_{\substack{i=1\\ y^{(i)}=0}}^{m} \nabla_{\mu_0} -\frac{1}{2}(x^{(i)}-\mu_0)^T\Sigma^{-1}(x^{(i)}-\mu_0).$$

$$= \sum_{\substack{i=1\\ y^i=0}}^{m} -\frac{1}{2}\cdot 2\Sigma^{-1}(x^{(i)}-\mu_0)(-1)$$

$$\sum_{\substack{i=1\\ y^i=0}}^{m} x^i - \sum_{\substack{i=1\\ y^i=0}}^{m} \mu_0 = 0$$

$$= \sum_{\substack{i=1\\ y^i=0}}^{m} \Sigma^{-1}(x^i-\mu_0) = 0. \Rightarrow \boxed{\sum_{\substack{i=1\\ y^i=0}}^{m}(x^i-\mu_0) = 0.}$$

$$\sum_{i=1}^{m} 1\{y^i=0\}x^i - \sum_{i=1}^{m} 1\{y^i=0\}\mu_0 = 0.$$

$$\mu_0 = \frac{\sum_{i=1}^{m} 1\{y^i=0\}x^{(i)}}{\sum_{i=1}^{m} 1\{y^i=0\}} \Bigg\} \text{mean of } x \text{ in class } 0.$$

$$\frac{\partial L}{\partial \mu_1} = 0 \Rightarrow \mu_1 = \frac{\sum_{i=1}^{m} 1\{y^i=1\}x^i}{\sum_{i=1}^{m} 1\{y^i=1\}}.$$

$$L = \sum_{i=1}^{m} \log \frac{1}{(2\pi)^{\frac{n}{2}}|\Sigma|^{\frac{1}{2}}} + \left(-\frac{1}{2}(x^{(i)}-\mu_{y^{(i)}})^T \Sigma^{-1}(x^{(i)}-\mu_{y^{(i)}})\right)$$

$$+ \sum_{i=1}^{m} y^i \log\phi + (1-y^i)\log(1-\phi)$$

$$= \sum_{i=1}^{m} \log\frac{1}{(2\pi)^{\frac{n}{2}}} + \sum_{i=1}^{m} \log|\Sigma|^{-\frac{1}{2}} - \frac{1}{2}(x^{(i)}-\mu_{y^{(i)}})^T\Sigma^{-1}(x^{(i)}-\mu_{y^{(i)}}) + \sum_{i=1}^{m} y^{(i)}\log\phi + (1-y^{(i)})\log(1-\phi)$$

② $\nabla_A |A| = |A|(A^{-1})^T$.

③ $\nabla_A\, x^T A^{-1} y$

$= -A^{-T} x y^T A^{-T}$.

$$\nabla_\Sigma L(\phi,\mu_0,\mu_1,\Sigma) = \nabla_\Sigma \sum_{i=1}^{m} \log|\Sigma|^{-\frac{1}{2}} - \frac{1}{2}(x^i-\mu_{y^i})^T\Sigma^{-1}(x^i-\mu_{y^i})$$

$$= \nabla_\Sigma \sum_{i=1}^{m} -\frac{1}{2}\log|\Sigma| - \frac{1}{2}(x^i-\mu_{y^i})^T\Sigma^{-1}(x^i-\mu_{y^i})$$

$$= \sum_{i=1}^{m} -\frac{1}{2}\frac{1}{\cancel{|\Sigma|}}\underbrace{\nabla_\Sigma|\Sigma|}_{\cancel{|\Sigma|}(\Sigma^{-1})^T} - \frac{1}{2}(-\Sigma^{-T})(x^i-\mu_{y^i})(x^i-\mu_{y^i})^T\Sigma^{-T}$$

~ symmetric

$$= \sum_{i=1}^{m} -\frac{1}{2}(\Sigma^{-1}) + \frac{1}{2}\Sigma^{-1}(x^i-\mu_{y^i})(x^i-\mu_{y^i})^T\Sigma^{-1} = 0.$$

$$\sum_{i=1}^{m}\left(1-\Sigma^{-1}(x^i-\mu_{y^i})(x^i-\mu_{y^i})^T\right) = 0.$$

$$m - \Sigma^{-1}\sum_{i=1}^{m}(x^i-\mu_{y^i})(x^i-\mu_{y^i})^T = 0.$$

$$\Sigma = \frac{1}{m}\sum_{i=1}^{m}(x^i-\mu_{y^i})(x^i-\mu_{y^i})^T.$$

Log likelihood of the data:

$$
\begin{aligned}
l(\phi, \mu_0, \mu_1, \Sigma) &= \log \prod_{i=1}^{m} p(x^{(i)}, y^{(i)}; \phi, \mu_0, \mu_1, \Sigma) \\
&= \log \prod_{i=1}^{m} p(x^{(i)} | y^{(i)}; \mu_0, \mu_1, \Sigma) p(y^{(i)}; \phi)
\end{aligned}
$$

Log likelihood of the data:

$$
\begin{aligned}
l(\phi, \mu_0, \mu_1, \Sigma) &= \log \prod_{i=1}^{m} p(x^{(i)}, y^{(i)}; \phi, \mu_0, \mu_1, \Sigma) \\
&= \log \prod_{i=1}^{m} p(x^{(i)} | y^{(i)}; \mu_0, \mu_1, \Sigma) p(y^{(i)}; \phi)
\end{aligned}
$$

Maximum likelihood estimate of the parameters:

$$
\left.
\begin{aligned}
\phi &= \frac{1}{m} \sum_{i=1}^{m} \mathbf{1}\{y^{(i)} = 1\} \\
\mu_b &= \frac{\sum_{i=1}^{m} \mathbf{1}\{y^{(i)} = b\} x^{(i)}}{\sum_{i=1}^{m} \mathbf{1}\{y^{(i)} = b\}} \text{ for } b = 0, 1 \\
\Sigma &= \frac{1}{m} \sum_{i=1}^{m} (x^{(i)} - \mu_{y^{(i)}})(x^{(i)} - \mu_{y^{(i)}})^T
\end{aligned}
\right]
$$

Maximum likelihood estimation of GDA

GDA finds a linear decision boundary at which
$p(y = 1|x) = p(y = 0|x) = 0.5$

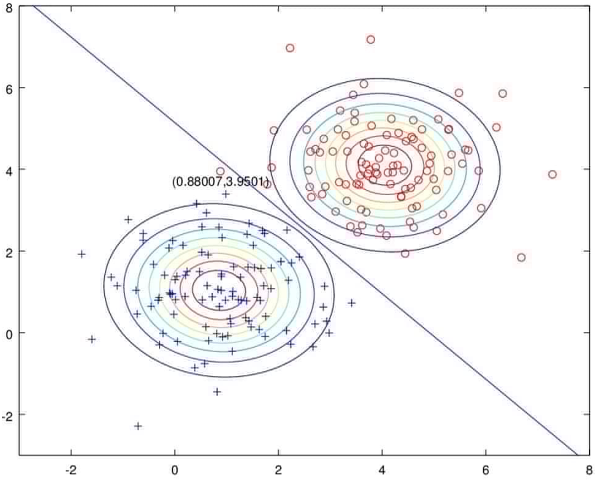

GDA and Logistic Regression

Proposition

$p(y=1|x;\phi,\mu_0,\mu_1,\Sigma)$ can be written in the form:

$$p(y=1|x;\phi,\Sigma,\mu_0,\mu_1) = \frac{1}{1+e^{-\theta^T x}}$$

GDA and Logistic Regression

Proposition

$p(y=1|x;\phi,\mu_0,\mu_1,\Sigma)$ can be written in the form:

$$p(y=1|x;\underbrace{\phi,\Sigma,\mu_0,\mu_1}_{H}) = \frac{1}{1+e^{-\theta^T x}}$$

$$\theta = \begin{bmatrix}\theta_1\\ \theta_2\end{bmatrix} \begin{matrix}\theta_0\\ = \\ \theta_1\end{matrix} \begin{bmatrix}\Sigma^{-1}(\mu_1-\mu_0)\\ \frac{1}{2}(\mu_0^T\Sigma^{-1}\mu_0 - \mu_1^T\Sigma^{-1}\mu_1) - \log\frac{1-\phi}{\phi}\end{bmatrix}, x = \begin{bmatrix}x_1\\ \vdots\\ x_n\\ 1\end{bmatrix}$$

$$p(y=1|x;H) = \frac{p(x|y=1;H)P(y=1;H)}{P(x;H)} = \frac{p(x|y=1;H)P(y=1;H)}{p(x|y=1;H)P(y=1;H) + p(x|y=0;H)P(y=0;H)}$$

$$= \frac{1}{1+\exp\big(\underbrace{-(\mu_1-\mu_0)\Sigma^{-1}}_{\theta_0}x - \underbrace{\big(\tfrac{1}{2}(\mu_0^T\Sigma^{-1}\mu_0-\mu_1^T\Sigma^{-1}\mu_1) - \log\tfrac{1-\phi}{\phi}}_{\theta_1 \in \mathbb{R}}\big)}$$

GDA and Logistic Regression

Proposition

$p(y=1|x;\phi,\mu_0,\mu_1,\Sigma)$ can be written in the form:

$$p(y=1|x;\phi,\Sigma,\mu_0,\mu_1) = \frac{1}{1+e^{-\theta^T x}}$$

$$\theta = \begin{bmatrix}\theta_1 \\ \theta_2\end{bmatrix} = \begin{bmatrix}\Sigma^{-1}(\mu_1-\mu_0) \\ \frac{1}{2}(\mu_0^T\Sigma^{-1}\mu_0 - \mu_1^T\Sigma^{-1}\mu_1) - \log\frac{1-\phi}{\phi}\end{bmatrix}, x = \begin{bmatrix}x_1 \\ \vdots \\ x_n \\ 1\end{bmatrix}$$

Similarly,

$$p(y=0|x;\phi,\Sigma,\mu_0,\mu_1) = \frac{1}{1+e^{\theta^T x}}$$

If $p(x|y) \sim \mathcal{N}(\mu,\Sigma)$, $p(y|x)$ is a logistic function.

GDA and Logistic Regression

$\Sigma_1 = \Sigma_2$

GDA

- Maximizes the **joint likelihood** $\prod_{i=1}^{m} p(x^{(i)}, y^{(i)})$
- Modeling assumptions: $x|y{=}b \sim \mathcal{N}(\mu_b, \Sigma)$, $y \sim \text{Bernoulli}(\phi)$ $P(y)$
- When modeling assumptions are correct, GDA is **asymptotically efficient** and **data efficient**

Logistic Regression

- Maximizes the **conditional likelihood** $\prod_{i=1}^{m} p(y^{(i)}|x^{(i)})$
- Modeling assumptions: $p(y|x)$ is a logistic function; no restriction on $p(x)$
- More robust and less sensitive to incorrect modeling assumptions.

Naïve Bayes

Naïve Bayes: Motivating Example

A simple generative learning algorithm for discrete input variables

Example: Spam filter (document classification)

Classify email messages x to spam $(y = 1)$ and non-spam $(y = 0)$ classes.

Hello [redacted]
We need to confirm your info...
(1) FINAL MESSAGE: Payout Verification - $3000 PAYOUT is ready to be addressed in your Name and we want to be sure it gets to the right place. Click below to start the confirmation process. The sooner you act, the sooner it can be in your hands!
Raging Bull Casino

A sample spam email

Example: Spam Filter

Binary text features

Given a dictionary of size n, represent a message composed of dictionary words as $x \in \{0, 1\}^n$:

$$x_i = \begin{cases} 1 & i\text{-th dictionary word is in message} \\ 0 & \text{otherwise} \end{cases}$$

$$x = \begin{array}{c|l} \begin{bmatrix} 0 \\ 0 \\ \vdots \\ 1 \\ \vdots \\ 1 \\ \vdots \\ 0 \end{bmatrix} & \begin{array}{l} \overset{A}{a} \\ aardvark \\ \vdots \\ casino \\ \vdots \\ payout \\ \vdots \\ \underset{z.}{zyzzyva} \end{array} \end{array} \Bigg\} n$$

Naïve Bayes Model

$p(y)$ $p(x|y)$

Probability of observing email $x_1, \ldots, x_n$ given spam class y :

$$p(x_1, \ldots, x_n|y) = p(x_1|y)p(x_2|y, x_1), \ldots, p(x_n|y, x_1, \ldots, x_{n-1})$$

Naïve Bayes (NB) assumption

x_i's are conditionally independent given y:

$$p(x_i|y, x_1, \ldots, x_{i-1}) = p(x_i|y)$$

size of dictionary

$$p(x|y) = p(x_1, \ldots, x_n|y) = p(x_1|y)p(x_2|y)\ldots p(x_n|y) = \prod_{i=1}^{n} p(x_i|y)$$

message

Naïve Bayes Parameters

Multi-variate Bernoulli event model

$x|y$ generated from n independent Bernoulli trials

$$p(x,y) = p(y)p(x|y) = p(y)\prod_{i=1}^{n} p(x_i|y)$$

- $y \sim Bernoulli(\phi_y)$: assume email class (spam vs no-spam) is randomly generated with prior $p(y) = \phi_y^y(1-\phi_y)^{1-y}$
- $x_i|y = b \sim Bernoulli(\phi_{i|y=b}),\ b = 1, 2$: given $y = b$, each word x_i is included in the message independently with $p(x_i = 1|y = b) = \phi_{i|y=b}$. i.e. (1, 0)

$p(x_i=1|y=1) = \phi_{i|y=1}$.
$p(x_i=0|y=0) = \phi_{i|y=0}$.

$$p(x_i|y=b) = \phi_{i|y=b}^{x_i}(1-\phi_{i|y=b})^{1-x_i}$$

Model parameters:

- ϕ_y } 1
- $\phi_{i|y=1}, \phi_{i|y=0}$ for $i = 1, \ldots, n$ 2n.

Naïve Bayes Parameter Learning

Likelihood of training data $(x^{(1)}, y^{(1)}), \ldots, (x^{(m)}, y^{(m)})$: i.i.d

$L(\theta)$

$$L(\phi_y, \phi_{j|y=0}, \phi_{j|y=1}) = \prod_{i=1}^{m} p(x^{(i)}, y^{(i)}) \quad \prod_{i=1}^{m} p(x^{(i)}|y^{(i)})\, p(y^{(i)})$$

$\log L(\cdots)$

Maximum likelihood estimation of parameters:

$$\phi_y = \frac{1}{m}\sum_{i=1}^{m} \mathbf{1}\{y^{(i)} = 1\}$$ *% of spam emails*

$y^{i} = 1$.

$\phi_1, \ldots, \phi_n$

$$\phi_{j|y=b} = \frac{\sum_{i=1}^{m} \mathbf{1}\{x_j^{(i)} = 1, y^{(i)} = b\}}{\sum_{i=1}^{m} \mathbf{1}\{y^{(i)} = b\}} \text{ for } b = 1, 0$$

if $b = 1$,

$$\frac{\#\{\text{ of messages of both spam and } x_j = 1\}}{\#\text{ spam messages}}$$

% of spam(non-spam) emails containing jth dictionary word

if $y^i = 0$.

$$\phi_{j|y=0} = \frac{\#\text{ of non spam messages containing } x_j}{\#\text{ of non spam messages}}$$

jth word appear in the ith message

=

Naïve Bayes Prediction

$[x_1, \ldots, x_n]$

Given new example with feature x, compute the posterior probability

known ϕ_y^*, $\phi_{i|y=b}$.

$$
\begin{aligned}
p(y=1|x) &= \frac{p(x|y=1)p(y=1)}{p(x)} \\
&= \frac{p(x|y=1)p(y=1)}{p(x|y=1)p(y=1)+p(x|y=0)p(y=0)} \\
&= \frac{\prod_{i=1}^{n} p(x_i|y=1)p(y=1)}{\prod_{i=1}^{n} p(x_i|y=1)p(y=1)+\prod_{i=1}^{n} p(x_i|y=0)p(y=0)}
\end{aligned}
$$

By the N.B assumption,

Choose label $y = 1$ (spam) if $p(y = 1|x) > T$ where $T \in [0, 1]$ is a threshold .. e.g. $T = 0.5$

T tradeoff between wrongly blocked non-spam (FPs) vs. wrongly blocked spams (FNs).

false negatives. false positives

Laplace smoothing

Issue with Naïve Bayes prediction:

- Suppose word x_j hasn't been seen in the training data,
  $\phi_{j|y=1} =$

For all $i = 1, \dots, m$, $1\{x_j^{(i)} = 1\} = 0$

$$\phi_{j|y=b} = \frac{\sum_{i=1}^{m} 1\{y^{i} = b, x_j^{(i)} = 1\}}{\sum_{i=1}^{m} 1\{y^{i} = b\}} \to 0.$$

$$= \qquad = 0.$$

$$\left.\begin{array}{l} \{\phi_{j|y=1}\} \\ \phi_{j|y=0} \end{array}\right\} = 0.$$

Given x, $\underline{P(y=1|x)} = \dfrac{\prod_{l=1}^{n} \underline{P(x_l|y=1)}_{0}\, \underline{P(y=1)}}{\prod_{l=1}^{n} \underline{P(x_l|y=1)} P(y=1) + \prod_{l=1}^{n} \underline{P(x_l|y=0)} P(y=0)}$ if $l = j$ $= \dfrac{0}{0}$

Laplace smoothing

Issue with Naïve Bayes prediction:

- Suppose word x_j hasn't been seen in the training data, $\phi_{j|y=1} = \phi_{j|y=0} = 0$
- Can not compute class posterior $p(y = 1|x) = \frac{0}{0}$.

Laplace smoothing

Issue with Naïve Bayes prediction:

- Suppose word x_j hasn't been seen in the training data, $\phi_{j|y=1} = \phi_{j|y=0} = 0$
- Can not compute class posterior $p(y = 1|x) = \frac{0}{0}$.

$y \in \{0, 1\}$ $k = 2$

$z \to y$

$z \to x|y$

$\phi_j = \frac{1}{m} \sum_{i=1}^{m} 1\{z^i = j\}$

Laplace smoothing

Let $z \in \{1, \ldots, k\}$ be a multinomial random variable. Given m independent observations $z^{(1)} \ldots z^{(m)}$, maximum likelihood estimation of $\phi_j = p(z = j)$ with **Laplace smoothing** is

$$\phi_j = \frac{\sum_{i=1}^{m} 1\{z^{(i)} = j\} + 1}{m + k}$$

$\sum_{j=1}^{n} \phi_j = 1$

ε

- $\phi_j \neq 0$ for all j
- $\sum_{j=1}^{k} \phi_j = 1$

Naïve Bayes with Laplace smoothing

Apply Laplace smoothing to $\phi_{j|y=b}$ for $b \in \{0, 1\}$

$$\phi_{j|y=b} = \frac{\sum_{i=1}^{m} \mathbf{1}\{x_j^{(i)} = 1, y^{(i)} = b\} + 1}{\sum_{i=1}^{m} \mathbf{1}\{y^i = b\} + 2}$$

In practice we don't apply Laplace smoothing to $\phi_y = p(y = 1)$, which is greater than 0.

Naïve Bayes Summary

Naïve Bayes (NB) assumption

x_i's are conditionally independent given y:

$$p(x_1, \ldots, x_n|y) = \prod_{i=1}^{n} p(x_i|y)$$

Different event models

- **Multi-variate Bernoulli model**: represent document of vocab size n as n independent Bernoulli trails
- **Multinomial event model**: represent document of N words as $x = \{x_1, \ldots, x_n\}$ where $x_i \in \{1, \ldots, K\}$. (not covered)

Homework

- Programming Assignment 1 late submission.
- TA sessions