Learning From Data
Lecture 4: Generative Learning Algorithms

Yang Li yangli@sz.tsinghua.edu.cn

October 9, 2020

Today's Lecture

Supervised Learning (Part II)

- Discriminative & Generative Models $x \in \mathbb{R}^n$
- Gaussian Discriminant Analysis ← binary classification
- Naïve Bayes $x \in \{1, \dots, k\}$ ← classification

  $x \in \{0,1\}^n$

Discriminative & Generative Models

Two Learning Approaches

blue red

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

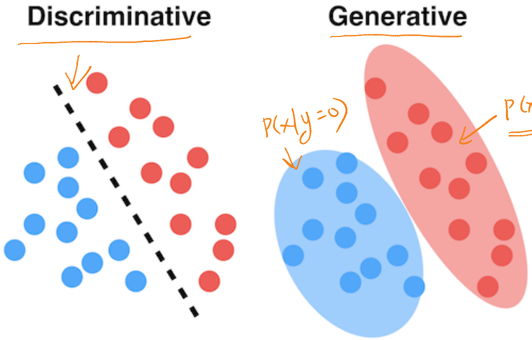

$p(x|y=0)$

$p(x|y=1)$

$p(y|x)$

Discriminate between classes of data points

Model the underlying distribution of the data

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 $\mathcal{Y}$.

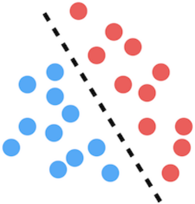

$p(y|x)$

- 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

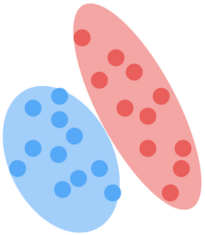

Bayes Rule

The posterior distribution on y given x:

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

Bayes Rule

The posterior distribution on y given x:

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

Make predictions in a generative model:

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

No need to calculate $p(x)$.

Generative Models

Generative classification algorithms:

- Continuous input: Gaussian Discriminant Analysis
- Discrete input: Naïve Bayes

generative

Gaussian Discriminant Analysis

Gaussian Discriminant Analysis: Overview

Linear Discriminant analysis (LDA) $\Sigma_1 = \Sigma_2$

Quadratic Discriminant analysis $\Sigma_1 \neq \Sigma_2$

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* .

$$y \sim Bernoulli(\phi) \quad P(y=1)=\phi \quad P(y=0)=1-\phi$$
$$x|y = 0 \sim N(\mu_0, \Sigma), x|y = 1 \sim N(\mu_1, \Sigma)$$

2. Estimate model parameters ϕ, μ_0 ,μ_1 and Σ from training data.
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 $N(\mu, \Sigma)$

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

Density function:

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

$\det(\Sigma)$

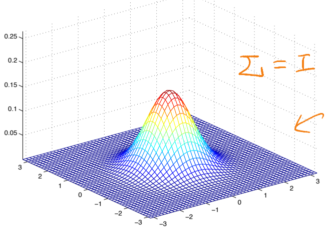

$\Sigma = I$

$x \in \mathbb{R}^2$

Multivariate Normal Distribution

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

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

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

Gaussian Discriminative Analysis

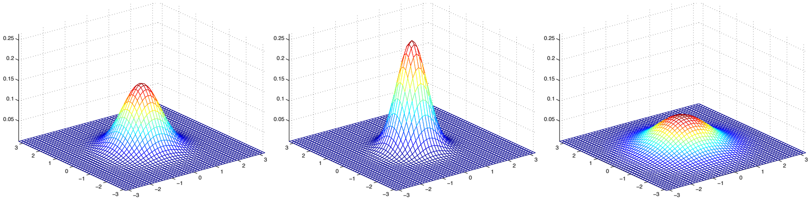

$$\Sigma = \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix} \qquad \Sigma = \begin{bmatrix} 0.6 & 0 \\ 0 & 0.6 \end{bmatrix} \qquad \Sigma = \begin{bmatrix} 2 & 0 \\ 0 & 2 \end{bmatrix}$$

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

Gaussian Discriminative Analysis

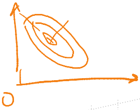

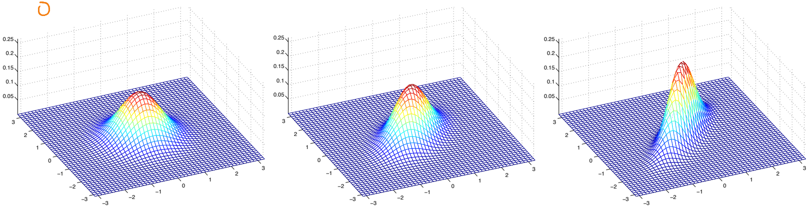

$$\Sigma = \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix} \qquad \Sigma = \begin{bmatrix} 1 & 0.5 \\ 0.5 & 1 \end{bmatrix} \qquad \Sigma = \begin{bmatrix} 1 & 0.8 \\ 0.8 & 1 \end{bmatrix}$$

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 $\phi, \mu_0, \mu_1, \Sigma$,

$$y \sim \text{Bernoulli}(\phi)$$
$$x|y = 0 \sim \mathcal{N}(\mu_0, \Sigma)$$
$$x|y = 1 \sim \mathcal{N}(\mu_1, \Sigma)$$

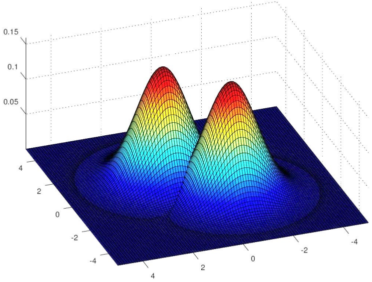

Probability density functions:

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

Log likelihood of the data: $(x^{(1)}, y^{(1)}) \dots (x^{(m)}, y^{(m)})$

$\in \mathbb{R}$ $\quad \in \mathbb{R}^n$ $\quad \in \mathbb{R}^{n\times n}$

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

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

$p(x^{(i)} | y^{(i)}; \mu_0, \mu_1, \Sigma) = \frac{1}{(2\pi)^{\frac{n}{2}} |\Sigma|^{\frac{1}{2}}} e^{-\frac{1}{2}(x-\mu_{y_i})^T \Sigma^{-1} (x-\mu_{y_i})}$

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

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

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

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

l

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

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

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

c : # of 1's in m samples

$m-c$: # of 0's in samples

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

$$\phi m - \phi c = c - c\phi$$

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

$$= \frac{1}{m}\sum_{i=1}^m 1\{y^i = 1\}$$

FACT ① $\nabla_A |A| = |A|(A^{-1})^T$ ② $\nabla_A x^TAx = xx^T$ ($tr(x^TAx)$) ③ $\nabla_A \ell(A) = 0 \Leftrightarrow \nabla_{A^{-1}} \ell(A) = 0.$ $|\Sigma|^{-1} = |\Sigma^{-1}|$

Log likelihood of the data:

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

$\nabla_S \log|S| = \frac{1}{|S|}\nabla_S |S| = \frac{1}{|S|} \cdot |S|(S^{-1})^T = (S^{-1})^T.$

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

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

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

$= \sum_{i=1, y^i=0}^{m} (\Sigma^{-1}x^i - \Sigma^{-1}\mu_0) = 0$

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

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

$\mu_0 = \frac{\sum_{i=1}^{m} 1\{y^i=0\}x^i}{\sum_{i=1}^{m} 1\{y^i=0\}}$ $\mu_1 = \frac{\sum_{i=1}^{m} 1\{y^i=1\}x^i}{\sum_{i=1}^{m} 1\{y^i=1\}}$

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

$|\Sigma^{-1}|^{\frac{1}{2}}$ $\frac{1}{2}\log|\Sigma^{-1}|$ $\boxed{S \triangleq \Sigma^{-1}}$

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

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

$\Rightarrow \sum_{i=1}^{m} \underbrace{(S^{-1})^T}_{\Sigma} = \sum_{i=1}^{m} (x^i - \mu_{y^i})(x^i - \mu_{y^i})^T$

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

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

Log likelihood of the data:

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

Maximum likelihood estimate of the parameters:

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

Maximum likelihood estimation of GDA

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

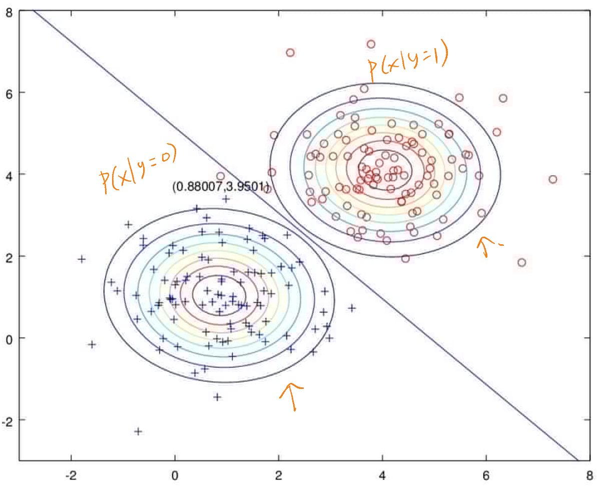

GDA and Logistic Regression

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

$$p(y = 1|x; \overbrace{\phi, \Sigma, \mu_0, \mu_1}^{H}) = \frac{1}{1 + e^{-\theta^T x}} + \theta_2$$

proof $P(y=1|x;H) = \frac{P(x|y=1;H)P(y=1;H)}{P(x;H)}$

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

$$= \frac{\frac{1}{A}\exp(-\frac{1}{2}(x-\mu_1)^T\Sigma^{-1}(x-\mu_1))\,\phi}{\frac{1}{A}\exp(-\frac{1}{2}(x-\mu_1)^T\Sigma^{-1}(x-\mu_1))\phi + \frac{1}{A}\exp(-\frac{1}{2}(x-\mu_0)^T\Sigma^{-1}(x-\mu_0))(1-\phi)}$$

$$= \frac{1}{1 + \frac{1-\phi}{\phi}\exp(-\frac{1}{2}(x-\mu_0)^T\Sigma^{-1}(x-\mu_0) + \frac{1}{2}(x-\mu_1)^T\Sigma^{-1}(x-\mu_1))}$$

$$= \frac{1}{1 + \exp((\mu_0-\mu_1)^T\Sigma^{-1}x + (\frac{1}{2}\mu_1^T\Sigma^{-1}\mu_1 - \frac{1}{2}\mu_0^T\Sigma^{-1}\mu_0 + \log\frac{1-\phi}{\phi}))}$$

$$= \frac{1}{1 + \exp(-(\underbrace{(\mu_1-\mu_0)\Sigma^{-1}}_{\in\mathbb{R}^n \;\; \theta_1}x + \underbrace{\frac{1}{2}(\mu_0^T\Sigma^{-1}\mu_0 - \mu_1^T\Sigma^{-1}\mu_1) - \log\frac{1-\phi}{\phi}}_{\in\mathbb{R} \;\; \theta_2}))}$$

GDA and Logistic Regression

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

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

where

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

GDA and Logistic Regression

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

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

where

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

$1-\frac{1}{1+e^{-\theta^T x}}$

$\Downarrow$

Similarly,

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

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

GDA and Logistic Regression

GDA

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

Logistic Regression

not considering $P(y)$

- 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: Motivatong 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 $x \in \{0, 1\}^n$:

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

$$x = \begin{array}{c} x_1 \\ x_2 \\ \vdots \\ x_i \\ \\ \\ \\ \\ \end{array} \begin{bmatrix} 0 \\ 0 \\ \vdots \\ 1 \\ \vdots \\ 1 \\ \vdots \\ 0 \end{bmatrix} \begin{array}{l} \textit{a} \\ \textit{aardvark} \\ \vdots \\ \textit{casino} \\ \vdots \\ \textit{payout} \\ \vdots \\ \textit{zyzzyva} \end{array}$$

Naïve Bayes Model

$$\begin{bmatrix} x_1 \\ \vdots \\ x_n \end{bmatrix}$$

$p(x \mid y=1)$

Probability of observing email $x_1, \ldots, x_n$ given spam class y : $= 1$

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

Naïve Bayes (NB) assumption

x_i's are conditionally independent given y:

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

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

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 $p(y) = \phi_y^y(1 - \phi_y)^{1-y}$ 1, 0
- $x_i|y = b \sim Bernoulli(\phi_{i|y=b}), b = \cancel{1, 2}$ 0 1 : 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.

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

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

Model parameters:

- ϕ_y
- $\phi_{i|y=1}, \phi_{i|y=0}$ for $i = 1, \ldots, n$ ← # of dictionary words

2n+1

Naïve Bayes Parameter Learning

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

for all j

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

Maximum likelihood estimation of parameters:

$$\phi_y = \frac{1}{m} \sum_{i=1}^{m} \mathbf{1}\{y^{(i)} = 1\}$$

% of spam emails

# of spam email/m

$p(x_j = 1 | y = b)$

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

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

Naïve Bayes Prediction

Given new example with feature x, compute the posterior probability

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

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

0.9

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

false positive

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

neg.

Laplace smoothing

$$\boxed{\phi_{j|y=b}} = \frac{\sum_{i=1}^{m} 1\{y^i = b, x_j^i = 1\}}{\sum_{i=1}^{m} 1\{y^i = b\}} \quad b = 1, 0$$

Issue with Naïve Bayes prediction:

- Suppose word x_j hasn't been seen in the training data,
  $\phi_{j|y=1} = 0 \;.=\; \phi_{j|y=0}$.

"and"

For all $i = 1, .., m$, $1\{y^i = b, x_j^i = 1\} = 0$

$$\Rightarrow \quad \phi_{j|y=1} = \frac{\sum_{i=1}^{m} 0}{\sum_{i=1}^{m} 1\{y^i = 1\}} = 0$$

similarly,

$$\phi_{j|y=0} = 0$$

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 $p(y = 1|x) = \frac{0}{0}$.

why? $\rightarrow$ $$P(y=1|x) = \frac{\prod_{l=1}^{n} P(X_l|y=1)P(y=1)}{\prod_{l=1}^{n} P(X_l|y=1)P(y=1) + \prod_{i=1}^{n} P(X_l|y=0)P(y=0)}$$

when $l=j$, $$P(X_l|y=1) = \phi_{l|y=1}^{x_l}(1-\phi_{l|y=1})^{1-x_l}$$
$$= \phi_{j|y=1}^{x_l}(1-\phi_{j|y=1})^{1-x_l}$$
$$= 0$$

Therefore $P(y=1|x) = \frac{0}{0+0}$

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 $p(y = 1|x) = \frac{0}{0}$.

x|y.

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

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

In Bernoulli Naive Bayes, $\chi = \{0, 1\}$, so $k = 2$

needed for $\sum_{j=1}^{k} \phi_j = 1$

make sure the "occurance" of every dictionary word is nonzero, to avoid numerical issues

"1" can be replaced by any small positive constant.

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

Naïve Bayes with Laplace smoothing

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

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

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

Naïve Bayes and Multinomial Event Model

Alternative text representation

- $x_i \in \{1, \ldots, K\}$ where K is the dictionary size
- Represent email of n words as $x = \{x_1, \ldots, x_n\}$ ← length of the message.

"a free gift..." $\rightarrow$ $\{x_1 = 1, x_2 = 1300, x_3 = 2433, \ldots\}$

dictionary id	1	2	...	1300	...	2433	...
word	a	aa	...	free	...	gift	...

$$x = \left.\begin{bmatrix} 1 \\ 1300 \\ 2433 \\ \vdots \end{bmatrix}\right\} n.$$

Naive Bayes and Multinomial Event Model

Multinomial event model

- first sampling $y \in \{0, 1\}$ from $p(y)$

$$y \sim Bernoulli(\phi_y)$$

- Select $x_1, x_2, \ldots, x_n$ independently from the same Multinomial distribution $p(x_i|y)$

$\phi_{k|y=b}$

$$x_i|y = b \sim Multinomial(\phi_{1|y=b}, \ldots, \phi_{K|y=b}), b = 0, 1$$

$$\boxed{\phi_{k|y=b}} = p(x_j = k|y = b) \text{ for all } j \in \{1, \ldots, n\}$$

kth word in the dictionary

For any word k in the dictionary, $\phi_{k|y}$ is the probability of k appear in an email given email class y

$P(X_1, \ldots, X_n | y)$

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

NB assumption

Multinomial event model parameters

position of word in message (does not matter)

Assume $p(x_j = k|y)$ is the same for all j

K : # of words in dictionary

- $\phi_y = p(y)$
- $\phi_{k|y=1} = p(x_j = k|y = 1)$ for $k = 1, \ldots, \cancel{n}\, K$
- $\phi_{k|y=0} = p(x_j = k|y = 0)$ for $k = 1, \ldots, \cancel{n}\, K$

$2K+1$ parameters

Likelihood of training set $(x^{(1)}, y^{(1)}), \ldots, (x^{(m)}, y^{(m)})$:

$$
\begin{aligned}
L(\phi_y, \phi_{k|y=0}, \phi_{k|y=1}) &= \prod_{i=1}^{m} p(x^{(i)}, y^{(i)}) \\
&= \prod_{i=1}^{m} p(x_1^{(i)}, \ldots, x_{n_i}^{(i)}, y^{(i)}) \\
&= \prod_{i=1}^{m} p(y^{i}; \phi_y)\, p(x_1, \ldots, x_{n_i} | y; \phi_{k|y=0}, \phi_{k|y=1}) \\
&= \prod_{i=1}^{m} p(y^{(i)}; \phi_y) \prod_{j=1}^{n_i} p(x_j^{(i)} | y; \phi_{k|y=0}, \phi_{k|y=1})
\end{aligned}
$$

length of the ith message

$$
K \left\{ \begin{bmatrix} \phi_{1|y=0} \\ \vdots \\ \phi_{K|y=0} \end{bmatrix} \right. \qquad \begin{bmatrix} \phi_{1|y=1} \\ \vdots \\ \phi_{K|y=1} \end{bmatrix}
$$

where n_i is the # words in the i-th email.

Maximum likelihood estimation with Laplace smoothing

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

$1\{x_j^i = 1, y^i = 1\}$.

- $\phi_{k|y=1} = \frac{\sum_{i=1}^{m} \left(\sum_{j=1}^{n_i} \mathbf{1}\{x_j^{(i)} = k, y^{(i)} = 1\}\right) + 1}{\sum_{i=1}^{m} \mathbf{1}\{y^{(i)} = 1\} n_i + K}$

# of times kth dictionary word appear in x^i given y^i is spam

- $\phi_{k|y=0} = \frac{\sum_{i=1}^{m} \sum_{j=1}^{n_i} \mathbf{1}\{x_j^{(i)} = k, y^{(i)} = 0\} + 1}{\sum_{i=1}^{m} \mathbf{1}\{y^{(i)} = 0\} n_i + K}$

K is the dictionary size.

① Bernoulli NB

② Multinomial NB.