

# Probabilistic Modeling

He He

CDS, NYU

March 24, 2020

# Contents

- 1 Overview
- 2 Conditional models
  - Recap: linear regression
  - A probabilistic view of linear regression
  - Logistic regression
  - Generalized Linear Models
- 3 Generative models
  - Naive Bayes models
    - Bernoulli Naive Bayes Models
    - Gaussian Naive Bayes Models

- Midterm (also see announcement on Brightspace)
  - Date and time: March 8 5pm–7pm ET
  - Location: remote on Gradescope; you will need to submit a pdf file or photos of your solution
  - Format: open book/notes but no internet
  - Coverage: up to kernel methods (not including this week)
  - Review: this week's lab
  - Difficulty: easier than last year
- Homeworks: the number of questions will be reduced for future homeworks

# Overview

# Why probabilistic modeling?

- A unified framework that covers many models, e.g., linear regression, logistic regression
- Learning as **statistical inference**
- Principled ways to incorporate your belief on the data generating distribution (inductive biases)

- Two ways to model how the data is generated:
  - **Conditional:**  $p(y | x)$
  - **Generative:**  $p(x, y)$
- How to estimate the parameters of our model? Maximum likelihood estimation.
- Compare and contrast conditional and generative models.

# Conditional models

# Linear regression

Linear regression is one of the most important methods in machine learning and statistics.

**Goal:** Predict a real-valued **target**  $y$  (also called response) from a vector of **features**  $x$  (also called covariates).

## Examples:

- Predicting house price given location, condition, build year etc.
- Predicting medical cost of a person given age, sex, region, BMI etc.
- Predicting age of a person based on their photos.



## Problem setup

**Data** Training examples  $\mathcal{D} = \{(x^{(n)}, y^{(n)})\}_{n=1}^N$ , where  $x \in \mathbb{R}^d$  and  $y \in \mathbb{R}$ .

**Model** A *linear* function  $h$  (parametrized by  $\theta$ ) to predict  $y$  from  $x$ :

$$h(x) = \sum_{i=0}^d \theta_i x_i = \theta^T x, \quad (1)$$

where  $\theta \in \mathbb{R}^d$  are the **parameters** (also called weights).

Note that

- We incorporate the **bias term** (also called the intercept term) into  $x$  (i.e.  $x_0 = 1$ ).
- We use superscript to denote the example id and subscript to denote the dimension id.

## Parameter estimation

**Loss function** We estimate  $\theta$  by minimizing the **squared loss** (the least square method):

$$J(\theta) = \frac{1}{N} \sum_{n=1}^N \left( y^{(n)} - \theta^T x^{(n)} \right)^2. \quad (\text{empirical risk}) \quad (2)$$

- Matrix form**
- Let  $X \in \mathbb{R}^{N \times d}$  be the **design matrix** whose rows are input features.
  - Let  $y \in \mathbb{R}^N$  be the vector of all targets.
  - We want to solve

$$\hat{\theta} = \arg \min_{\theta} (X\theta - y)^T (X\theta - y). \quad (3)$$

**Solution** Closed-form solution:  $\hat{\theta} = (X^T X)^{-1} X^T y$ .

### Review questions

- Derive the solution for linear regression.
- What if  $X^T X$  is not invertible?

We've seen

- Linear regression: response is a linear function of the inputs
- Estimate parameters by minimize the squared loss

But...

- Why squared loss is a reasonable choice for regression problems?
- What assumptions are we making on the data? (**inductive bias**)

Next,

- Derive linear regression from a **probabilistic modeling perspective**.

# Assumptions in linear regression

- $x$  and  $y$  are related through a linear function:

$$y = \theta^T x + \epsilon, \quad (4)$$

where  $\epsilon$  is the **residual error** capturing all unmodeled effects (e.g., noise).

- The errors are distributed *iid* (independently and identically distributed):

$$\epsilon \sim \mathcal{N}(0, \sigma^2). \quad (5)$$

What's the distribution of  $Y \mid X = x$ ?

$$p(y \mid x; \theta) = \mathcal{N}(\theta^T x, \sigma^2). \quad (6)$$

Imagine putting a Gaussian bump around the output of the linear predictor.

# Maximum likelihood estimation (MLE)

Given a probabilistic model and a dataset  $\mathcal{D}$ , how to estimate the model parameters  $\theta$ ?

The **maximum likelihood principle** says that we should maximize the (conditional) likelihood of the data:

$$L(\theta) \stackrel{\text{def}}{=} p(\mathcal{D}; \theta) \tag{7}$$

$$= \prod_{n=1}^N p(y^{(n)} | x^{(n)}; \theta). \quad (\text{examples are distributed } iid) \tag{8}$$

In practice, we maximize the **log likelihood**  $\ell(\theta)$ , or equivalently, minimize the negative log likelihood (NLL).

## MLE for linear regression

Let's find the MLE solution for our model. Recall that  $Y | X = x \sim \mathcal{N}(\theta^T x, \sigma^2)$ .

$$\ell(\theta) \stackrel{\text{def}}{=} \log L(\theta) \tag{9}$$

$$= \log \prod_{n=1}^N p(y^{(n)} | x^{(n)}; \theta) \tag{10}$$

$$= \sum_{n=1}^N \log p(y^{(n)} | x^{(n)}; \theta) \tag{11}$$

$$= \sum_{n=1}^N \log \frac{1}{\sqrt{2\pi}\sigma} \exp\left(-\frac{(y^{(n)} - \theta^T x^{(n)})^2}{2\sigma^2}\right) \tag{12}$$

$$= N \log \frac{1}{\sqrt{2\pi}\sigma} - \frac{1}{2\sigma^2} \sum_{n=1}^N (y^{(n)} - \theta^T x^{(n)})^2 \tag{13}$$

## Gradient of the likelihood

Recall that we obtained the normal equation by setting the derivative of the squared loss to zero. Now let's compute the derivative of the likelihood w.r.t. the parameters.

$$\ell(\theta) = N \log \frac{1}{\sqrt{2\pi\sigma}} - \frac{1}{2\sigma^2} \sum_{n=1}^N \left( y^{(n)} - \theta^T x^{(n)} \right)^2 \quad (14)$$

$$\frac{\partial \ell}{\partial \theta_i} = -\frac{1}{\sigma^2} \sum_{n=1}^N (y^{(n)} - \theta^T x^{(n)}) x_i^{(n)}. \quad (15)$$

(Spoiler: we will see this form again.)

We've seen

- Linear regression assumes that  $Y | X = x$  follows a Gaussian distribution
- MLE of linear regression is equivalent to the least square method

However,

- Sometimes Gaussian distribution is not a reasonable assumption, e.g., classification
- Can we use the same modeling approach for other prediction tasks?

Next,

- Derive [logistic regression](#) for classification.



## Assumptions in logistic regression

Consider binary classification where  $Y \in \{0, 1\}$ . What should be the distribution  $Y | X = x$ ?

We model  $p(y | x)$  as a **Bernoulli** distribution:

$$p(y | x) = h(x)^y (1 - h(x))^{1-y}. \quad (16)$$


How should we parameterize  $h(x)$ ?

- What is  $p(y = 1 | x)$  and  $p(y = 0 | x)$ ?  $h(x) \in (0, 1)$ .
- What is the mean of  $Y | X = x$ ?  $h(x)$ . (Think how we parameterize the mean in linear regression)
- Need a function  $f$  to map the linear predictor  $\theta^T x$  in  $\mathbb{R}$  to  $(0, 1)$ :

$$f(\eta) = \frac{1}{1 + e^{-\eta}} \quad \text{logistic function} \quad (17)$$

# Logistic regression

$$f(\eta) = \frac{1}{1+e^{-\eta}}$$



- $p(y | x) = \text{Bernoulli}(f(\theta^T x))$ .
- When do we have  $p(y = 1 | x) = 1$  and  $p(y = 0 | x) = 1$ ?
- **Exercise:** show that the **log odds** is

$$\log \frac{p(y = 1 | x)}{p(y = 0 | x)} = \theta^T x. \quad (18)$$

$$\implies \text{linear decision boundary} \quad (19)$$

- How do we extend it to multiclass classification? (more on this later)

## MLE for logistic regression

Similar to linear regression, let's estimate  $\theta$  by maximizing the conditional log likelihood.

$$\ell(\theta) = \sum_{n=1}^N \log p(y^{(n)} | x^{(n)}; \theta) \quad (20)$$

$$= \sum_{n=1}^N y^{(n)} \log f(\theta^T x^{(n)}) + (1 - y^{(n)}) \log(1 - f(\theta^T x^{(n)})) \quad (21)$$

- Closed-form solutions are not available.
- But, the likelihood is concave—**gradient ascent** gives us the unique optimal solution.

$$\theta := \theta + \alpha \nabla_{\theta} \ell(\theta). \quad (22)$$

# Gradient descent for logistic regression

## Math review: Chain rule

If  $z$  depends on  $y$  which itself depends on  $x$ , e.g.,  $z = (y(x))^2$ , then  $\frac{dz}{dx} = \frac{dz}{dy} \frac{dy}{dx}$ .

Likelihood for a single example:  $\ell^n = y^{(n)} \log f(\theta^T x^{(n)}) + (1 - y^{(n)}) \log(1 - f(\theta^T x^{(n)}))$ .

$$\frac{\partial \ell^n}{\partial \theta_i} = \frac{\partial \ell^n}{\partial f^n} \frac{\partial f^n}{\partial \theta_i} \quad (23)$$

$$= \left( \frac{y^{(n)}}{f^n} - \frac{1 - y^{(n)}}{1 - f^n} \right) \frac{\partial f^n}{\partial \theta_i} \quad \frac{d}{dx} \ln x = \frac{1}{x} \quad (24)$$

$$= \left( \frac{y^{(n)}}{f^n} - \frac{1 - y^{(n)}}{1 - f^n} \right) \left( f^n (1 - f^n) x_i^{(n)} \right) \quad \text{Exercise: apply chain rule to } \frac{\partial f^n}{\partial \theta_i} \quad (25)$$

$$= (y^{(n)} - f^n) x_i^{(n)} \quad \text{simplify by algebra} \quad (26)$$

The full gradient is thus  $\frac{\partial \ell}{\partial \theta_i} = \sum_{n=1}^N (y^{(n)} - f(\theta^T x^{(n)})) x_i^{(n)}$ .

## A closer look at the gradient

$$\frac{\partial \ell}{\partial \theta_i} = \sum_{n=1}^N (y^{(n)} - f(\theta^T x^{(n)})) x_i^{(n)} \quad (27)$$

- Does this look familiar?
- Our derivation for linear regression and logistic regression are quite similar...
- Next, a more general family of models.

## Compare linear regression and logistic regression

	linear regression	logistic regression
Combine the inputs	$\theta^T x$ (linear)	$\theta^T x$ (linear)
Output	real	categorical
Conditional distribution	Gaussian	Bernoulli
Transfer function $f(\theta^T x)$	identity	logistic
Mean $\mathbb{E}(Y   X = x; \theta)$	$f(\theta^T x)$	$f(\theta^T x)$

- $x$  enters through a linear function.
- The main **difference** between the formulations is due to different conditional distributions.
- Can we generalize the idea to handle other output types, e.g., positive integers?

# Construct a generalized regression model

**Task:** Given  $x$ , predict  $p(y | x)$

**Modeling:**

- Choose a parametric family of distributions  $p(y; \theta)$  with parameters  $\theta \in \Theta$
- Choose a transfer function that maps a linear predictor in  $\mathbb{R}$  to  $\Theta$

$$\underbrace{x}_{\in \mathbb{R}^d} \mapsto \underbrace{w^T x}_{\in \mathbb{R}} \mapsto \underbrace{f(w^T x)}_{\in \Theta} = \theta, \quad (28)$$

**Learning:** MLE:  $\hat{\theta} \in \arg \max_{\theta} \log p(\mathcal{D}; \hat{\theta})$

**Inference:** For prediction, use  $x \rightarrow f(w^T x)$

## Example: Construct Poisson regression

Say we want to predict the number of people entering a restaurant in New York during lunch time.

- What features would be useful?
- What's a good model for number of visitors (the **output distribution**)?

### Math review: Poisson distribution

Given a random variable  $Y \in 0, 1, 2, \dots$  following  $\text{Poisson}(\lambda)$ , we have

$$p(Y = k; \lambda) = \frac{\lambda^k e^{-\lambda}}{k!}, \quad (29)$$

where  $\lambda > 0$  and  $\mathbb{E}[Y] = \lambda$ .

The Poisson distribution is usually used to model the number of events occurring during a fixed period of time.



## Example: Construct Poisson regression

We've decided that  $Y | X = x \sim \text{Poisson}(\eta)$ , what should be the transfer function  $f$ ?  
 $x$  enters linearly:

$$x \mapsto \underbrace{w^T x}_R \mapsto \lambda = \underbrace{f(w^T x)}_{(0, \infty)}$$

Standard approach is to take

$$f(w^T x) = \exp(w^T x).$$

Likelihood of the full dataset  $\mathcal{D} = \{(x_1, y_1), \dots, (x_n, y_n)\}$ :

$$\log p(y_i; \lambda_i) = [y_i \log \lambda_i - \lambda_i - \log(y_i!)] \quad (30)$$

$$\log p(\mathcal{D}; w) = \sum_{i=1}^n [y_i \log [\exp(w^T x_i)] - \exp(w^T x_i) - \log(y_i!)] \quad (31)$$

$$= \sum_{i=1}^n [y_i w^T x_i - \exp(w^T x_i) - \log(y_i!)] \quad (32)$$

## Example: multinomial logistic regression

How to extend logistic regression to multiclass classification?

Output: Bernoulli distribution  $\rightarrow$  **categorical distribution**

- Parametrized by a probability vector  $\theta = (\theta_1, \dots, \theta_k) \in \mathbb{R}^k$ :
  - $\sum_{i=1}^k \theta_i = 1$  and  $\theta_i \geq 0$  for  $i = 1, \dots, k$
  - So  $\forall y \in \{1, \dots, k\}$ ,  $p(y) = \theta_y$ .

From each  $x$ , we compute a linear score function for each class:

$$x \mapsto (\langle w_1, x \rangle, \dots, \langle w_k, x \rangle) \in \mathbb{R}^k,$$

What's the transfer function that maps this  $\mathbb{R}^k$  vector into a probability?

The **softmax function**:

$$(s_1, \dots, s_k) \mapsto \theta = \left( \frac{e^{s_1}}{\sum_{i=1}^k e^{s_i}}, \dots, \frac{e^{s_k}}{\sum_{i=1}^k e^{s_i}} \right).$$

# Multinomial Logistic Regression

- Say we want to get the predicted categorical distribution for a given  $x \in \mathbb{R}^d$ .
- First compute the scores ( $\in \mathbb{R}^k$ ) and then their softmax:

$$x \mapsto (\langle w_1, x \rangle, \dots, \langle w_k, x \rangle) \mapsto \theta = \left( \frac{\exp(w_1^T x)}{\sum_{i=1}^k \exp(w_i^T x)}, \dots, \frac{\exp(w_k^T x)}{\sum_{i=1}^k \exp(w_i^T x)} \right)$$

- We can write the conditional probability for any  $y \in \{1, \dots, k\}$  as

$$p(y | x; w) = \frac{\exp(w_y^T x)}{\sum_{i=1}^k \exp(w_i^T x)}.$$

Recipe for constructing a conditional distribution for prediction:

- 1 Define input and output space (as for any other model).
- 2 Choose the output distribution  $p(y | x; \theta)$  based on the task
- 3 Choose the transfer function that maps  $w^T x$  to a  $\Theta$ .
- 4 (The formal family is called “generalized linear models”.)

Learning:

- Fit the model by maximum likelihood estimation.
- Closed solutions do not exist in general, so we use gradient ascent.

# Generative models

We've seen

- Model the conditional distribution  $p(y | x; \theta)$  using generalized linear models.
- (Previously) Directly map  $x$  to  $y$ , e.g., perceptron.

Next,

- Model the **joint distribution**  $p(x, y; \theta)$ .
- Predict the label for  $x$  as  $\arg \max_{y \in \mathcal{Y}} p(x, y; \theta)$ .

# Generative modeling through the Bayes rule

Training:

$$p(x, y) = p(x | y)p(y) \quad (33)$$

Testing:

$$p(y | x) = \frac{p(x | y)p(y)}{p(x)} \quad \text{Bayes rule} \quad (34)$$

$$\arg \max_y p(y | x) = \arg \max_y p(x | y)p(y) \quad (35)$$

## Naive Bayes (NB) models

Let's consider binary text classification (e.g., fake vs genuine review) as a motivating example.

**Bag-of-words** representation of a document

- ["machine", "learning", "is", "fun", "."]
- $x_i \in \{0, 1\}$ : whether the  $i$ -th word in our vocabulary exists in the input

$$x = [x_1, x_2, \dots, x_d] \quad \text{where } d = \text{vocabulary size} \quad (36)$$

What's the probability of a document  $x$ ?

$$p(x | y) = p(x_1, \dots, x_d | y) \quad (37)$$

$$= p(x_1 | y) p(x_2 | y, x_1) p(x_3 | y, x_2, x_1) \dots p(x_d | y, x_{d-1}, \dots, x_1) \quad \text{chain rule} \quad (38)$$

$$= \prod_{i=1}^d p(x_i | y, x_{<i}) \quad (39)$$



## Naive Bayes assumption

**Challenge:**  $p(x_i | y, x_{<i})$  is hard to model (and estimate), especially for large  $i$ .

Solution:

### Naive Bayes assumption

Features are **conditionally independent** given the label:

$$p(x | y) = \prod_{i=1}^d p(x_i | y). \quad (40)$$

A strong assumption in general, but works well in practice.

## Parametrize $p(x_i | y)$ and $p(y)$

For binary  $x_i$ , assume  $p(x_i | y)$  follows Bernoulli distributions.

$$p(x_i = 1 | y = 1) = \theta_{i,1}, \quad p(x_i = 1 | y = 0) = \theta_{i,0}. \quad (41)$$

Similarly,

$$p(y = 1) = \theta_0. \quad (42)$$

Thus,

$$p(x, y) = p(x | y)p(y) \quad (43)$$

$$= p(y) \prod_{i=1}^d p(x_i | y) \quad \text{NB assumption} \quad (44)$$

$$= p(y) \prod_{i=1}^d \theta_{i,y} \mathbb{I}\{x_i = 1\} + (1 - \theta_{i,y}) \mathbb{I}\{x_i = 0\} \quad (45)$$

Indicator function  $\mathbb{I}\{\text{condition}\}$  evaluates to 1 if “condition” is true and 0 otherwise.

## MLE for our NB model

We maximize the likelihood of the data  $\prod_{n=1}^N p_{\theta}(x^{(n)}, y^{(n)})$  (as opposed to the *conditional* likelihood we've seen before).

$$\frac{\partial}{\partial \theta_{j,1}} \ell = \frac{\partial}{\partial \theta_{j,1}} \sum_{n=1}^N \sum_{i=1}^d \log \left( \theta_{i,y^{(n)}} \mathbb{I} \{x_i^{(n)} = 1\} + (1 - \theta_{i,y^{(n)}}) \mathbb{I} \{x_i^{(n)} = 0\} \right) + \log p_{\theta_0}(y^{(n)}) \quad (46)$$

$$= \frac{\partial}{\partial \theta_{j,1}} \sum_{n=1}^N \log \left( \theta_{j,y^{(n)}} \mathbb{I} \{x_j^{(n)} = 1\} + (1 - \theta_{j,y^{(n)}}) \mathbb{I} \{x_j^{(n)} = 0\} \right) \quad \text{ignore } i \neq j \quad (47)$$

$$= \sum_{n=1}^N \mathbb{I} \{y^{(n)} = 1 \wedge x_j^{(n)} = 1\} \frac{1}{\theta_{j,1}} + \mathbb{I} \{y^{(n)} = 1 \wedge x_j^{(n)} = 0\} \frac{1}{1 - \theta_{j,1}} \quad \text{ignore } y^{(n)} = 0 \quad (48)$$

## MLE solution for our NB model

Set  $\frac{\partial}{\partial \theta_{j,1}} \ell$  to zero:

$$\theta_{j,1} = \frac{\sum_{n=1}^N \mathbb{I}\{y^{(n)} = 1 \wedge x_j^{(n)} = 1\}}{\sum_{n=1}^N \mathbb{I}\{y^{(n)} = 1\}} \quad (49)$$

In practice, count words:

$$\frac{\text{number of fake reviews containing "absolutely"}}{\text{number of fake reviews}}$$

**Exercise:** show that

$$\theta_{j,0} = \frac{\sum_{n=1}^N \mathbb{I}\{y^{(n)} = 0 \wedge x_j^{(n)} = 1\}}{\sum_{n=1}^N \mathbb{I}\{y^{(n)} = 0\}} \quad (50)$$

$$\theta_0 = \frac{\sum_{n=1}^N \mathbb{I}\{y^{(n)} = 1\}}{N} \quad (51)$$

NB assumption: **conditionally independent** features given the label

Recipe for learning a NB model:

- 1 Choose  $p(x_i | y)$ , e.g., Bernoulli distribution for binary  $x_i$ .
- 2 Choose  $p(y)$ , often a categorical distribution.
- 3 Estimate parameters by MLE (same as the strategy for conditional models) .

Next, NB with continuous features.

## NB with continuous inputs

Let's consider a multiclass classification task with continuous inputs.

$$p(x_i | y) \sim \mathcal{N}(\mu_{i,y}, \sigma_{i,y}^2) \quad (52)$$

$$p(y = k) = \theta_k \quad (53)$$

Likelihood of the data:

$$p(\mathcal{D}) = \prod_{n=1}^N p(y^{(n)}) \prod_{i=1}^d p(x_i^{(n)} | y^{(n)}) \quad (54)$$

$$= \prod_{n=1}^N \theta_{y^{(n)}} \prod_{i=1}^d \frac{1}{\sqrt{2\pi}\sigma_{i,y^{(n)}}} \exp\left(-\frac{1}{2\sigma_{i,y^{(n)}}^2} (x_i^{(n)} - \mu_{i,y^{(n)}})^2\right) \quad (55)$$

# MLE for Gaussian NB

Log likelihood:

$$\ell = \sum_{n=1}^N \log \theta_{y^{(n)}} + \sum_{n=1}^N \sum_{i=1}^d \log \frac{1}{\sqrt{2\pi}\sigma_{i,y^{(n)}}} - \frac{1}{2\sigma_{i,y^{(n)}}^2} \left( x_i^{(n)} - \mu_{i,y^{(n)}} \right)^2 \quad (56)$$

$$\frac{\partial}{\partial \mu_{j,k}} \ell = \frac{\partial}{\partial \mu_{j,k}} \sum_{n:y^{(n)}=k} -\frac{1}{2\sigma_{j,k}^2} \left( x_j^{(n)} - \mu_{j,k} \right)^2 \quad \text{ignore irrelevant terms} \quad (57)$$

$$= \sum_{n:y^{(n)}=k} \frac{1}{\sigma_{j,k}^2} \left( x_j^{(n)} - \mu_{j,k} \right) \quad (58)$$

Set  $\frac{\partial}{\partial \mu_{j,k}} \ell$  to zero:

$$\mu_{j,k} = \frac{\sum_{n:y^{(n)}=k} x_j^{(n)}}{\sum_{n:y^{(n)}=k} 1} = \text{sample mean of } x_j \text{ in class } k \quad (59)$$

Exercise: show that

$$\sigma_{j,k}^2 = \frac{\sum_{n:y^{(n)}=k} (x_j^{(n)} - \mu_{j,k})^2}{\sum_{n:y^{(n)}=k} 1} = \text{sample variance of } x_j \text{ in class } k \quad (60)$$

$$\theta_k = \frac{\sum_{n:y^{(n)}=k} 1}{N} \quad (\text{class prior}) \quad (61)$$



## Decision boundary of the Gaussian NB model

Is the Gaussian NB model a linear classifier?

$$\log \frac{p(y = 1 | x)}{p(y = 0 | x)} = \log \frac{p(x | y = 1)p(y = 1)}{p(x | y = 0)p(y = 0)} \quad (62)$$

$$= \log \frac{\theta_0}{1 - \theta_0} + \sum_{i=1}^d \left( \log \sqrt{\frac{\sigma_{i,0}^2}{\sigma_{i,1}^2}} + \left( \frac{(x_i - \mu_{i,0})^2}{2\sigma_{i,0}^2} - \frac{(x_i - \mu_{i,1})^2}{2\sigma_{i,1}^2} \right) \right) \quad \text{quadratic} \quad (63)$$

$$\text{assume that } \sigma_{i,0} = \sigma_{i,1} = \sigma_i, \quad (\theta_0 = 0.5) \quad (64)$$

$$= \sum_{i=1}^d \frac{1}{2\sigma_i^2} \left( (x_i - \mu_{i,0})^2 - (x_i - \mu_{i,1})^2 \right) \quad (65)$$

$$= \sum_{i=1}^d \frac{\mu_{i,1} - \mu_{i,0}}{\sigma_i^2} x_i + \frac{\mu_{i,0}^2 - \mu_{i,1}^2}{2\sigma_i^2} \quad \text{linear} \quad (66)$$

## Decision boundary of the Gaussian NB model

Assuming the variance of each feature is the same for both classes, we have

$$\log \frac{p(y = 1 | x)}{p(y = 0 | x)} = \sum_{i=1}^d \frac{\mu_{i,1} - \mu_{i,0}}{\sigma_i^2} x_i + \frac{\mu_{i,0}^2 - \mu_{i,1}^2}{2\sigma_i^2} \quad (67)$$

$$= \theta^T x \quad \text{where else have we seen it?} \quad (68)$$

$$(69)$$

$$\theta_i = \frac{\mu_{i,1} - \mu_{i,0}}{\sigma_i^2} \quad \text{for } i \in [1, d] \quad (70)$$

$$\theta_0 = \sum_{i=1}^d \frac{\mu_{i,0}^2 - \mu_{i,1}^2}{2\sigma_i^2} \quad \text{bias term} \quad (71)$$

## Naive Bayes vs logistic regression

	logistic regression	Gaussian naive Bayes
model type	conditional/discriminative	generative
parametrization	$p(y   x)$	$p(x   y), p(y)$
assumptions on $Y$	Bernoulli	Bernoulli
assumptions on $X$	—	Gaussian
decision boundary	$\theta_{LR}^T x$	$\theta_{GNB}^T x$

Given the same training data, is  $\theta_{LR} = \theta_{GNB}$ ?


# Naive Bayes vs logistic regression

Logistic regression and Gaussian naive Bayes converge to the same classifier asymptotically, assuming the GNB assumption holds.

What if the GNB assumption is not true?

## Generative vs discriminative classifiers

Ng, A. and Jordan, M. (2002). [On discriminative versus generative classifiers: A comparison of logistic regression and naive Bayes](#). In Advances in Neural Information Processing Systems 14.



Solid line: naive Bayes; dashed line: logistic regression.