



Generative Models




Discriminative vs Generative Models



Build a classifier that learns a decision boundary


Classify test data by seeing which side of the boundary it falls into. (what are some examples we've seen?)



Build a model for what the data in each class looks like (i.e., the distribution of the data)


Classify test data by comparing it against the models for the data in the two classes.

Discriminative vs Generative Models



Build a classifier that
learns a decision boundary


Models $p(y|x)$ directly



Build a model of what data
in each class looks like


Models $p(x|y)$ for each
value of y and $p(y)$ and
then use Bayes' rule to find
 $p(y|x)$

Example of $p(x|y)$



- $P(x_1, x_2 | y = \text{blue}) = N(x_1 | \mu = 1, \sigma^2 = 1) N(x_2 | \mu = 1, \sigma^2 = 1)$
 - Product of normal densities for both x_1 and x_2
 - Highest density at (1, 1), lower densities far away from (1,1)
- Wouldn't work for the red crosses: the probability distribution doesn't seem symmetric

Example of $p(y)$



- $P(y = \text{blue}) = \frac{12}{12+14}, P(y = \text{red}) = \frac{14}{12+14}$

Does the patient have cancer?

- We perform a lab test and the result comes back positive
- The test comes back positive in 98% of cases where the cancer is present
- The test comes back negative in 97% of cases where there is no cancer
- 0.008 of the population has cancer
 - And the cancer screening was random
 - (Why is this important?)

Generative process for data

- $P(\text{cancer}) = 0.008$
- $P(+|\text{cancer}) = .98$
- $P(-|\neg\text{cancer}) = .97$

θ

- θ determines how the test results are generated
 - Person i has cancer with prob. 0.008
 - The probability of a positive test for person i depends on whether they have cancer or not

- $P(\text{cancer}) = 0.008$
- $P(+|\text{cancer}) = .98$
- $P(-|\neg\text{cancer}) = .97$

- $$P(\text{cancer}|+) = \frac{P(+|\text{cancer})P(\text{cancer})}{P(+)}$$

$$= \frac{P(+|\text{cancer})P(\text{cancer})}{P(+|\text{cancer})P(\text{cancer}) + P(+|\neg\text{cancer})P(\neg\text{cancer})}$$

Learning a Generative Model

- For the cancer data, just count the number of points in the training set (of size N) belonging to each category
- $P(\text{cancer}) \approx \frac{\text{count}(\text{cancer},)}{N}$
- $P(+|\text{cancer}) \approx \frac{\text{count}(\text{cancer},+)}{\text{count}(\text{cancer})}$
- (Could get $P(\text{cancer}|+)$ by counting as well)

Gaussian Classifiers

- Suppose the test actually outputs a real number t
 - $P(\text{cancer}) \approx \frac{\text{count}(\text{cancer})}{N}$
 - $P(t|\text{cancer}) = N(t|\mu_{\text{cancer}}, \sigma_{\text{cancer}}^2)$
 - $P(t|\neg\text{cancer}) = N(t|\mu_{\neg\text{cancer}}, \sigma_{\neg\text{cancer}}^2)$
 - $\theta = \{\mu_{\text{cancer}}, \mu_{\neg\text{cancer}}, \sigma_{\text{cancer}}, \sigma_{\neg\text{cancer}}, \dots\}$
 - θ determines how the test results are generated
 - Decide whether person i has cancer (with prob $P(\text{cancer})$)
 - Now generate the test output t
- What's the probability that the person has cancer, if we know θ ?
 - $$P_{\theta}(\text{cancer}|t) = \frac{P_{\theta}(t|\text{cancer})P_{\theta}(\text{cancer})}{P_{\theta}(t|\text{cancer})P_{\theta}(\text{cancer}) + P_{\theta}(t|\neg\text{cancer})P_{\theta}(\neg\text{cancer})}$$

Learn using
maximum
likelihood

Learning a Gaussian with Maximum Likelihood

- We have all the t 's for patients *with* cancer
- Maximum Likelihood:
 - $\operatorname{argmax}_{\mu_c, \sigma_c^2} \prod_i N(t^{(i)} | \mu_c, \sigma_c^2), i \in \text{cancer}$
 - Solution (show with calculus!):
 - $\widehat{\mu}_c = \overline{t^{(i)}}, i \in \text{cancer}$
 - $\widehat{\sigma}_c^2 = \sum_i \frac{(t^{(i)} - \widehat{\mu}_c)^2}{\#\text{ncancer}}, i \in \text{cancer}$

Classification of new instances

- Suppose don't know θ
- What's $P(\text{cancer}|t, D)$?
 - *Not* $P_{\theta_{MAP}}(\text{cancer}|t, D)$!

Classification of new instances

- Suppose we are estimating θ from the data
- What's $P(\text{cancer}|t, D)$?
 - $\sum_{\theta' \in \Theta} P(\text{cancer}|\theta', t)P(\theta'|D) = \sum_{\theta' \in \Theta} P_{\theta'}(\text{cancer}|t)P(\theta'|D)$
 - Intuition: consider all the possible θ' , compute the probability according to each of them, and weight them by how much we believe that the true θ could be θ'

- Suppose we are estimating θ from the data
- What's $P(\text{cancer}|t)$?
 - $P_{\theta_{MAP}}(\text{cancer}|t)$ is not a horrible estimate here