
Text Categorization

1

Categorization

- **Given:**
 - A description of an instance, $x \in X$, where X is the *instance language* or *instance space*.
 - A fixed set of categories:
 $C = \{c_1, c_2, \dots, c_n\}$
- **Determine:**
 - The category of x : $c(x) \in C$, where $c(x)$ is a categorization function whose domain is X and whose range is C .

2

Learning for Categorization

- A training example is an instance $x \in X$, paired with its correct category $c(x)$: $\langle x, c(x) \rangle$ for an unknown categorization function, c .
- Given a set of training examples, D .
- Find a hypothesized categorization function, $h(x)$, such that:
$$\forall \langle x, c(x) \rangle \in D : h(x) = c(x)$$

Consistency

3

Sample Category Learning Problem

- Instance language: $\langle \text{size, color, shape} \rangle$
 - size $\in \{\text{small, medium, large}\}$
 - color $\in \{\text{red, blue, green}\}$
 - shape $\in \{\text{square, circle, triangle}\}$
- $C = \{\text{positive, negative}\}$

• D :

Example	Size	Color	Shape	Category
1	small	red	circle	positive
2	large	red	circle	positive
3	small	red	triangle	negative
4	large	blue	circle	negative

4

General Learning Issues

- Many hypotheses are usually consistent with the training data.
- Bias
 - Any criteria other than consistency with the training data that is used to select a hypothesis.
- Classification accuracy (% of instances classified correctly).
 - Measured on independent test data.
- Training time (efficiency of training algorithm).
- Testing time (efficiency of subsequent classification).

5

Generalization

- Hypotheses must generalize to correctly classify instances not in the training data.
- Simply memorizing training examples is a consistent hypothesis that does not generalize.
- *Occam's razor*:
 - Finding a *simple* hypothesis helps ensure generalization.

6

Text Categorization

- Assigning documents to a fixed set of categories.
- Applications:
 - Web pages
 - Recommending
 - Yahoo-like classification
 - Newsgroup Messages
 - Recommending
 - spam filtering
 - News articles
 - Personalized newspaper
 - Email messages
 - Routing
 - Prioritizing
 - Folderizing
 - spam filtering

7

Learning for Text Categorization

- Manual development of text categorization functions is difficult.
- Learning Algorithms:
 - Bayesian (naïve)
 - Neural network
 - Relevance Feedback (Rocchio)
 - Rule based (Ripper)
 - Nearest Neighbor (case based)
 - Support Vector Machines (SVM)


8

Using Relevance Feedback (Rocchio)

- Relevance feedback methods can be adapted for text categorization.
- Use standard TF/IDF weighted vectors to represent text documents (normalized by maximum term frequency).
- For each category, compute a *prototype* vector by summing the vectors of the training documents in the category.
- Assign test documents to the category with the closest prototype vector based on cosine similarity.

9

Illustration of Rocchio Text Categorization



10

Rocchio Text Categorization Algorithm (Training)

Assume the set of categories is $\{c_1, c_2, \dots, c_n\}$
For i from 1 to n let $\mathbf{p}_i = \langle 0, 0, \dots, 0 \rangle$ (*init. prototype vectors*)
For each training example $\langle x, c(x) \rangle \in D$
Let \mathbf{d} be the frequency normalized TF/IDF term vector for doc x
Let $i = j: (c_j = c(x))$
(*sum all the document vectors in c_i to get \mathbf{p}_i*)
Let $\mathbf{p}_i = \mathbf{p}_i + \mathbf{d}$

11

Rocchio Text Categorization Algorithm (Test)

Given test document x
Let \mathbf{d} be the TF/IDF weighted term vector for x
Let $m = -2$ (*init. maximum cosSim*)
For i from 1 to n :
(*compute similarity to prototype vector*)
Let $s = \text{cosSim}(\mathbf{d}, \mathbf{p}_i)$
if $s > m$
let $m = s$
let $r = c_i$ (*update most similar class prototype*)
Return class r

12

Rocchio Properties

- Does not guarantee a consistent hypothesis.
- Forms a simple generalization of the examples in each class (a *prototype*).
- Prototype vector does not need to be averaged or otherwise normalized for length since cosine similarity is insensitive to vector length.
- Classification is based on similarity to class prototypes.

13

Rocchio Time Complexity

- **Note:** The time to add two sparse vectors is proportional to minimum number of non-zero entries in the two vectors.
- **Training Time:** $O(|D|(L_d + |V_d|)) = O(|D| L_d)$ where L_d is the average length of a document in D and $|V_d|$ is the average vocabulary size for a document in D .
- **Test Time:** $O(L_t + |C||V_t|)$ where L_t is the average length of a test document and $|V_t|$ is the average vocabulary size for a test document.
 - Assumes lengths of \mathbf{p}_i vectors are computed and stored during training, allowing $\text{cosSim}(\mathbf{d}, \mathbf{p}_i)$ to be computed in time proportional to the number of non-zero entries in \mathbf{d} (i.e. $|V_t|$)

14

Nearest-Neighbor Learning Algorithm

- Learning is just storing the representations of the training examples in D .
- Testing instance x :
 - Compute similarity between x and all examples in D .
 - Assign x the category of the most similar example in D .
- Does not explicitly compute a generalization or category prototypes.
- Also called:
 - Case-based
 - Memory-based
 - Lazy learning

15

K Nearest-Neighbor

- Using only the closest example to determine categorization is subject to errors due to:
 - A single atypical example.
 - Noise (i.e. error) in the category label of a single training example.
- More robust alternative is to find the k most-similar examples and return the majority category of these k examples.
- Value of k is typically odd to avoid ties, 3 and 5 are most common.


16

Similarity Metrics

- Nearest neighbor method depends on a similarity (or distance) metric.
- Simplest for continuous m -dimensional instance space is *Euclidian distance*.
- Simplest for m -dimensional binary instance space is *Hamming distance* (number of feature values that differ).
- For text, cosine similarity of TF-IDF weighted vectors is typically most effective.


17

3 Nearest Neighbor Illustration (Euclidian Distance)



18

Illustration of 3 Nearest Neighbor for Text



19

K Nearest Neighbor for Text

Training:

For each training example $\langle x, c(x) \rangle \in D$
Compute the corresponding TF-IDF vector, \mathbf{d}_x , for document x

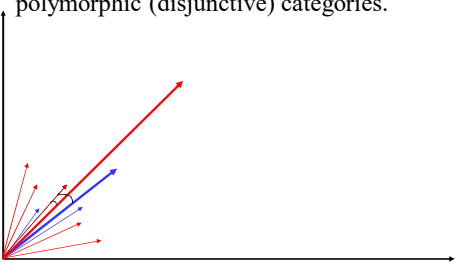
Test instance y :

Compute TF-IDF vector \mathbf{d} for document y
For each $\langle x, c(x) \rangle \in D$
Let $s_x = \text{cosSim}(\mathbf{d}, \mathbf{d}_x)$
Sort examples, x , in D by decreasing value of s_x
Let N be the first k examples in D . (*get most similar neighbors*)
Return the majority class of examples in N

20

Rocchio Anomaly


- Prototype models have problems with polymorphic (disjunctive) categories.



21

3 Nearest Neighbor Comparison

- Nearest Neighbor tends to handle polymorphic categories better.



22

Nearest Neighbor Time Complexity

- **Training Time:** $O(|D| L_d)$ to compose TF-IDF vectors.
- **Testing Time:** $O(L_t + |D| |V_t|)$ to compare to all training vectors.
 - Assumes lengths of \mathbf{d}_t vectors are computed and stored during training, allowing $\text{cosSim}(\mathbf{d}, \mathbf{d}_t)$ to be computed in time proportional to the number of non-zero entries in \mathbf{d} (i.e. $|V_t|$)
- Testing time can be high for large training sets.

23

Nearest Neighbor with Inverted Index

- Determining k nearest neighbors is the same as determining the k best retrievals using the test document as a query to a database of training documents.
- Use standard VSR inverted index methods to find the k nearest neighbors.
- **Testing Time:** $O(B |V_t|)$ where B is the average number of training documents in which a test-document word appears.
- Therefore, overall classification is $O(L_t + B |V_t|)$
 - Typically $B \ll |D|$

24


Bayesian Methods

- Learning and classification methods based on probability theory.
- Bayes theorem plays a critical role in probabilistic learning and classification.
- Uses *prior* probability of each category given no information about an item.
- Categorization produces a *posterior* probability distribution over the possible categories given a description of an item.

25

Axioms of Probability Theory

- All probabilities between 0 and 1
 $0 \leq P(A) \leq 1$
- True proposition has probability 1, false has probability 0.
 $P(\text{true}) = 1 \quad P(\text{false}) = 0.$
- The probability of disjunction is:
 $P(A \vee B) = P(A) + P(B) - P(A \wedge B)$




26

Conditional Probability

- $P(A | B)$ is the probability of A given B
- Assumes that B is all and only information known.
- Defined by:

$$P(A | B) = \frac{P(A \wedge B)}{P(B)}$$



27

Independence

- A and B are independent iff:

$$P(A|B) = P(A) \quad \text{These two constraints are logically equivalent}$$

$$P(B|A) = P(B)$$

- Therefore, if A and B are independent:

$$P(A|B) = \frac{P(A \wedge B)}{P(B)} = P(A)$$

$$P(A \wedge B) = P(A)P(B)$$

28

Joint Distribution

- The joint probability distribution for a set of random variables, X_1, \dots, X_n gives the probability of every combination of values (an n -dimensional array with v^n values if all variables are discrete with v values, all v^n values must sum to 1): $P(X_1, \dots, X_n)$

	positive		negative	
	circle	square	circle	square
red	0.20	0.02	0.05	0.30
blue	0.02	0.01	0.20	0.20

- The probability of all possible conjunctions (assignments of values to some subset of variables) can be calculated by summing the appropriate subset of values from the joint distribution.

$$P(\text{red} \wedge \text{circle}) = 0.20 + 0.05 = 0.25$$

$$P(\text{red}) = 0.20 + 0.02 + 0.05 + 0.30 = 0.57$$

- Therefore, all conditional probabilities can also be calculated.

$$P(\text{positive} | \text{red} \wedge \text{circle}) = \frac{P(\text{positive} \wedge \text{red} \wedge \text{circle})}{P(\text{red} \wedge \text{circle})} = \frac{0.20}{0.25} = 0.80$$

29

Probabilistic Classification

- Let Y be the random variable for the class which takes values $\{y_1, y_2, \dots, y_m\}$.
- Let X be the random variable describing an instance consisting of a vector of values for n features $\langle X_1, X_2, \dots, X_n \rangle$, let x_i be a possible value for X and x_{ij} a possible value for X_i .
- For classification, we need to compute $P(Y=y_i | X=x_i)$ for $i=1 \dots m$.
- However, given no other assumptions, this requires a table giving the probability of each category for each possible instance in the instance space, which is impossible to accurately estimate from a reasonably-sized training set.
 - Assuming Y and all X_i are binary, we need 2^n entries to specify $P(Y=\text{pos} | X=x_i)$ for each of the 2^n possible x_i 's since $P(Y=\text{neg} | X=x_i) = 1 - P(Y=\text{pos} | X=x_i)$
 - Compared to $2^{n+1} - 1$ entries for the joint distribution $P(Y, X_1, X_2, \dots, X_n)$

30

Bayes Theorem

$$P(H | E) = \frac{P(E | H)P(H)}{P(E)}$$

Simple proof from definition of conditional probability:

$$P(H | E) = \frac{P(H \wedge E)}{P(E)} \quad (\text{Def. cond. prob.})$$

$$P(E | H) = \frac{P(H \wedge E)}{P(H)} \quad (\text{Def. cond. prob.})$$

$$P(H \wedge E) = P(E | H)P(H)$$

QED: $P(H | E) = \frac{P(E | H)P(H)}{P(E)}$

31

Bayesian Categorization

- Determine category of x_k by determining for each y_i

$$P(Y = y_i | X = x_k) = \frac{P(Y = y_i)P(X = x_k | Y = y_i)}{P(X = x_k)}$$

- $P(X = x_k)$ can be determined since categories are complete and disjoint.

$$\sum_{i=1}^m P(Y = y_i | X = x_k) = \sum_{i=1}^m \frac{P(Y = y_i)P(X = x_k | Y = y_i)}{P(X = x_k)} = 1$$

$$P(X = x_k) = \sum_{i=1}^m P(Y = y_i)P(X = x_k | Y = y_i)$$

32

Bayesian Categorization (cont.)

- Need to know:
 - Priors: $P(Y = y_i)$
 - Conditionals: $P(X = x_k | Y = y_i)$
- $P(Y = y_i)$ are easily estimated from data.
 - If n_i of the examples in D are in y_i then $P(Y = y_i) = n_i / |D|$
- Too many possible instances (e.g. 2^n for binary features) to estimate all $P(X = x_k | Y = y_i)$.
- Still need to make some sort of independence assumptions about the features to make learning tractable.

33

Generative Probabilistic Models


- Assume a simple (usually unrealistic) probabilistic method by which the data was generated.
- For categorization, each category has a different parameterized generative model that characterizes that category.
- Training:** Use the data for each category to estimate the parameters of the generative model for that category.
 - Maximum Likelihood Estimation (MLE):** Set parameters to maximize the probability that the model produced the given training data.
 - If M_λ denotes a model with parameter values λ and D_k is the training data for the k th class, find model parameters for class k (λ_k) that maximize the likelihood of D_k :

$$\lambda_k = \underset{\lambda}{\operatorname{argmax}} P(D_k | M_\lambda)$$


- Testing:** Use Bayesian analysis to determine the category model that most likely generated a specific test instance.

34

Naïve Bayes Generative Model





Positive




Negative

35

Naïve Bayes Inference Problem



Positive



Negative

36

Naïve Bayesian Categorization

- If we assume features of an instance are independent **given the category** (*conditionally independent*).

$$P(X | Y) = P(X_1, X_2, \dots, X_n | Y) = \prod_{i=1}^n P(X_i | Y)$$

- Therefore, we then only need to know $P(X_i | Y)$ for each possible pair of a feature-value and a category.
- If Y and all X_i are binary, this requires specifying only $2n$ parameters:
 - $P(X_i = \text{true} | Y = \text{true})$ and $P(X_i = \text{true} | Y = \text{false})$ for each X_i
 - $P(X_i = \text{false} | Y) = 1 - P(X_i = \text{true} | Y)$
- Compared to specifying 2^n parameters without any independence assumptions.

37

Naïve Bayes Example

Probability	positive	negative
$P(Y)$	0.5	0.5
$P(\text{small} Y)$	0.4	0.4
$P(\text{medium} Y)$	0.1	0.2
$P(\text{large} Y)$	0.5	0.4
$P(\text{red} Y)$	0.9	0.3
$P(\text{blue} Y)$	0.05	0.3
$P(\text{green} Y)$	0.05	0.4
$P(\text{square} Y)$	0.05	0.4
$P(\text{triangle} Y)$	0.05	0.3
$P(\text{circle} Y)$	0.9	0.3

Test Instance:
<medium, red, circle>

38

Naïve Bayes Example

Probability	positive	negative
$P(Y)$	0.5	0.5
$P(\text{medium} Y)$	0.1	0.2
$P(\text{red} Y)$	0.9	0.3
$P(\text{circle} Y)$	0.9	0.3

Test Instance:
<medium, red, circle>

$$P(\text{positive} | X) = \frac{P(\text{positive}) \cdot P(\text{medium} | \text{positive}) \cdot P(\text{red} | \text{positive}) \cdot P(\text{circle} | \text{positive})}{P(X)}$$

$$= \frac{0.5 \cdot 0.1 \cdot 0.9 \cdot 0.9}{0.0405 / P(X)} = 0.0405 / 0.0495 = 0.8181$$

$$P(\text{negative} | X) = \frac{P(\text{negative}) \cdot P(\text{medium} | \text{negative}) \cdot P(\text{red} | \text{negative}) \cdot P(\text{circle} | \text{negative})}{P(X)}$$

$$= \frac{0.5 \cdot 0.2 \cdot 0.3 \cdot 0.3}{0.009 / P(X)} = 0.009 / 0.0495 = 0.1818$$

$$P(\text{positive} | X) + P(\text{negative} | X) = 0.0405 / P(X) + 0.009 / P(X) = 1$$

$$P(X) = (0.0405 + 0.009) = 0.0495$$

39

Estimating Probabilities

- Normally, probabilities are estimated based on observed frequencies in the training data.
- If D contains n_k examples in category y_k , and n_{ijk} of these n_k examples have the j th value for feature X_j , x_{ij} , then:

$$P(X_i = x_{ij} | Y = y_k) = \frac{n_{ijk}}{n_k}$$

- However, estimating such probabilities from small training sets is error-prone.
- If due only to chance, a rare feature, X_j , is always false in the training data, $\forall y_k : P(X_j = \text{true} | Y = y_k) = 0$.
- If $X_j = \text{true}$ then occurs in a test example, X , the result is that $\forall y_k : P(X | Y = y_k) = 0$ and $\forall y_k : P(Y = y_k | X) = 0$

40

Probability Estimation Example

Ex	Size	Color	Shape	Category	Probability	positive	negative
					P(Y)		
1	small	red	circle	positive	P(small Y)	0.5	0.5
					P(medium Y)	0.0	0.0
2	large	red	circle	positive	P(large Y)	0.5	0.5
					P(red Y)	1.0	0.5
3	small	red	triangle	negative	P(blue Y)	0.0	0.5
					P(green Y)	0.0	0.0
4	large	blue	circle	negative	P(square Y)	0.0	0.0
					P(triangle Y)	0.0	0.5
					P(circle Y)	1.0	0.5

Test Instance X :
<medium, red, circle>

$P(\text{positive} | X) = 0.5 * 0.0 * 1.0 * 1.0 / P(X) = 0$
 $P(\text{negative} | X) = 0.5 * 0.0 * 0.5 * 0.5 / P(X) = 0$

41

Smoothing

- To account for estimation from small samples, probability estimates are adjusted or *smoothed*.
- Laplace smoothing using an m -estimate assumes that each feature is given a prior probability, p , that is assumed to have been previously observed in a “virtual” sample of size m .

$$P(X_i = x_{ij} | Y = y_k) = \frac{n_{ijk} + mp}{n_k + m}$$

- For binary features, p is simply assumed to be 0.5.

42

Laplace Smoothing Example

- Assume training set contains 10 positive examples:
 - 4: small
 - 0: medium
 - 6: large
- Estimate parameters as follows (if $m=1, p=1/3$)
 - $P(\text{small} \mid \text{positive}) = (4 + 1/3) / (10 + 1) = 0.394$
 - $P(\text{medium} \mid \text{positive}) = (0 + 1/3) / (10 + 1) = 0.03$
 - $P(\text{large} \mid \text{positive}) = (6 + 1/3) / (10 + 1) = 0.576$
 - $P(\text{small or medium or large} \mid \text{positive}) = 1.0$

43

Naïve Bayes for Text

- Modeled as generating a bag of words for a document in a given category by repeatedly sampling with replacement from a vocabulary $V = \{w_1, w_2, \dots, w_m\}$ based on the probabilities $P(w_j \mid c_i)$.
- Smooth probability estimates with Laplace m -estimates assuming a uniform distribution over all words ($p = 1/|V|$) and $m = |V|$
 - Equivalent to a virtual sample of seeing each word in each category exactly once.

44

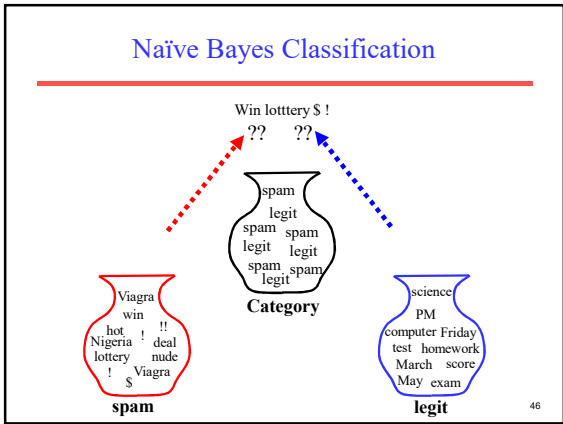
Naïve Bayes Generative Model for Text

spam

Category

legit

45



Text Naïve Bayes Algorithm (Train)

Let V be the vocabulary of all words in the documents in D

For each category $c_i \in C$

- Let D_i be the subset of documents in D in category c_i
- $P(c_i) = |D_i| / |D|$
- Let T_i be the concatenation of all the documents in D_i
- Let n_i be the total number of word occurrences in T_i
- For each word $w_j \in V$
 - Let n_{ij} be the number of occurrences of w_j in T_i
 - Let $P(w_j | c_i) = (n_{ij} + 1) / (n_i + |V|)$

47

Text Naïve Bayes Algorithm (Test)

Given a test document X

Let n be the number of word occurrences in X

Return the category:

$$\operatorname{argmax}_{c_i \in C} P(c_i) \prod_{i=1}^n P(a_i | c_i)$$

where a_i is the word occurring the i th position in X

48

Underflow Prevention

- Multiplying lots of probabilities, which are between 0 and 1 by definition, can result in floating-point underflow.
- Since $\log(xy) = \log(x) + \log(y)$, it is better to perform all computations by summing logs of probabilities rather than multiplying probabilities.
- Class with highest final un-normalized log probability score is still the most probable.

49

Naïve Bayes Posterior Probabilities

- Classification results of naïve Bayes (the class with maximum posterior probability) are usually fairly accurate.
- However, due to the inadequacy of the conditional independence assumption, the actual posterior-probability numerical estimates are not.
 - Output probabilities are generally very close to 0 or 1.

50

Evaluating Categorization

- Evaluation must be done on test data that are independent of the training data (usually a disjoint set of instances).
- *Classification accuracy*: c/n where n is the total number of test instances and c is the number of test instances correctly classified by the system.
- Results can vary based on sampling error due to different training and test sets.
- Average results over multiple training and test sets (splits of the overall data) for the best results.

51

N-Fold Cross-Validation

- Ideally, test and training sets are independent on each trial.
 - But this would require too much labeled data.
- Partition data into N equal-sized disjoint segments.
- Run N trials, each time using a different segment of the data for testing, and training on the remaining $N-1$ segments.
- This way, at least test-sets are independent.
- Report average classification accuracy over the N trials.
- Typically, $N = 10$.

52

Learning Curves

- In practice, labeled data is usually rare and expensive.
- Would like to know how performance varies with the number of training instances.
- *Learning curves* plot classification accuracy on independent test data (Y axis) versus number of training examples (X axis).

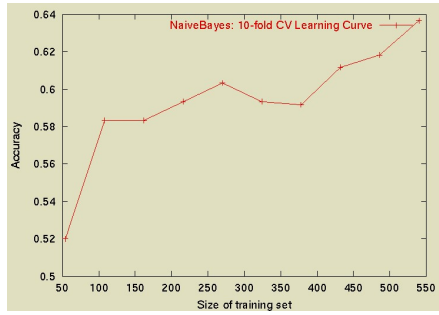
53

N-Fold Learning Curves

- Want learning curves averaged over multiple trials.
- Use N -fold cross validation to generate N full training and test sets.
- For each trial, train on increasing fractions of the training set, measuring accuracy on the test data for each point on the desired learning curve.

54

Sample Learning Curve (Yahoo Science Data)



55
