

CS540 Introduction to Artificial Intelligence

Lecture 7

Young Wu

Based on lecture slides by Jerry Zhu, Yingyu Liang, and Charles Dyer

June 20, 2022

K Nearest Neighbor

Description

- Given a new instance, find the K instances in the training set that are the closest.
- Predict the label of the new instance by the majority of the labels of the K instances.

Distance Function

Definition

- Many distance functions can be used in place of the Euclidean distance.

$$\rho(x, x') = \|x - x'\|_2 = \sqrt{\sum_{j=1}^m (x_j - x'_j)^2}$$

- An example is Manhattan distance.

$$\rho(x, x') = \sum_{j=1}^m |x_j - x'_j|$$

p Norms

Definition

- Another group of examples is the p norms.

$$\rho(x, x') = \left(\sum_{j=1}^m |x_j - x'_j|^p \right)^{\frac{1}{p}}$$

- $p = 1$ is the Manhattan distance.
- $p = 2$ is the Euclidean distance.
- $p = \infty$ is the sup distance, $\rho(x, x') = \max_{i=1,2,\dots,m} \{|x_j - x'_j|\}$.
- p cannot be less than 1.

K Nearest Neighbor Algorithm

- Input: instances: $\{x_i\}_{i=1}^n$ and $\{y_i\}_{i=1}^n$, and a new instance \hat{x} .
- Output: new label \hat{y} .
- Order the training instances according to the distance to \hat{x} .

$$\rho(\hat{x}, x_{(i)}) \leq \rho(\hat{x}, x_{(i+1)}), i = 1, 2, \dots, n - 1$$

- Assign the majority label of the closest k instances.

$$\hat{y} = \text{mode} \{y_{(1)}, y_{(2)}, \dots, y_{(k)}\}$$

Zipf's Law

Motivation

- If the word count is f and the word rank is r , then
$$f \cdot r \approx \text{constant}$$
- This relation is called Zipf's Law

Bag of Words Features

Definition

- Given a document i and vocabulary with size m , let c_{ij} be the count of the word j in the document i for $j = 1, 2, \dots, m$.
- Bag of words representation of a document has features that are the count of each word divided by the total number of words in the document.

$$x_{ij} = \frac{c_{ij}}{\sum_{j'=1}^m c_{ij'}}$$

TF IDF Features

Definition

- Another feature representation is called tf-idf, which stands for normalized term frequency, inverse document frequency.

$$\text{tf}_{ij} = \frac{c_{ij}}{\max_{j'} c_{ij'}}, \quad \text{idf}_j = \log \frac{n}{\sum_{i=1}^n \mathbb{1}_{\{c_{ij} > 0\}}}$$

$$x_{ij} = \text{tf}_{ij} \text{idf}_j$$

- n is the total number of documents and $\sum_{i=1}^n \mathbb{1}_{\{c_{ij} > 0\}}$ is the number of documents containing word j .

Cosine Similarity

Definition

- The similarity of two documents i and i' is often measured by the cosine of the angle between the feature vectors.

$$\text{sim}(x_i, x_{i'}) = \frac{x_i^T x_{i'}}{\sqrt{x_i^T x_i} \sqrt{(x_{i'})^T x_{i'}}}$$

N -Gram Model

Description

- Count all n gram occurrences.
- Apply Laplace smoothing to the counts.
- Compute the conditional transition probabilities.

Token Notations

Definition

- A word (or character) at position t of a sentence (or string) is denoted as z_t .
- A sentence (or string) with length d is (z_1, z_2, \dots, z_d) .
- $\mathbb{P}\{Z_t = z_t\}$ is the probability of observing $z_t \in \{1, 2, \dots, j\}$ at position t of the sentence, usually shortened to $\mathbb{P}\{z_t\}$.

Unigram Model

Definition

- Unigram models assume independence.

$$\mathbb{P}\{z_1, z_2, \dots, z_d\} = \prod_{t=1}^d \mathbb{P}\{z_t\}$$

- In general, two events A and B are independent if:

$$\mathbb{P}\{A|B\} = \mathbb{P}\{A\} \text{ or } \mathbb{P}\{A, B\} = \mathbb{P}\{A\} \mathbb{P}\{B\}$$

- For a sequence of words, independence means:

$$\mathbb{P}\{z_t|z_{t-1}, z_{t-2}, \dots, z_1\} = \mathbb{P}\{z_t\}$$

Maximum Likelihood Estimation

Definition

- $\mathbb{P}\{z_t\}$ can be estimated by the count of the word z_t .

$$\hat{\mathbb{P}}\{z_t\} = \frac{c_{z_t}}{\sum_{z=1}^m c_z}$$

- This is called the maximum likelihood estimator because it maximizes the probability of observing the sentences in the training set.

MLE Example

Definition

- Let $p = \hat{\mathbb{P}}\{0\}$ in a string with c_0 0's and c_1 1's.
- The probability of observing the string is:

$$\binom{c_0 + c_1}{c_0} p^{c_0} (1 - p)^{c_1}$$

- The above expression is maximized by:

$$p^* = \frac{c_0}{c_0 + c_1}$$

Bigram Model

Definition

- Bigram models assume Markov property.

$$\mathbb{P}\{z_1, z_2, \dots, z_d\} = \mathbb{P}\{z_1\} \prod_{t=2}^d \mathbb{P}\{z_t | z_{t-1}\}$$

- Markov property means the distribution of an element in the sequence only depends on the previous element.

$$\mathbb{P}\{z_t | z_{t-1}, z_{t-2}, \dots, z_1\} = \mathbb{P}\{z_t | z_{t-1}\}$$

Conditional Probability

Definition

- In general, the conditional probability of an event A given another event B is the probability of A and B occurring at the same time divided by the probability of event B .

$$\mathbb{P}\{A|B\} = \frac{\mathbb{P}\{AB\}}{\mathbb{P}\{B\}}$$

- For a sequence of words, the conditional probability of observing z_t given z_{t-1} is observed is the probability of observing both divided by the probability of observing z_{t-1} first.

$$\mathbb{P}\{z_t|z_{t-1}\} = \frac{\mathbb{P}\{z_{t-1}, z_t\}}{\mathbb{P}\{z_{t-1}\}}$$

Bigram Model Estimation

Definition

- Using the conditional probability formula, $\mathbb{P}\{z_t|z_{t-1}\}$, called transition probabilities, can be estimated by counting all bigrams and unigrams.

$$\hat{\mathbb{P}}\{z_t|z_{t-1}\} = \frac{c_{z_{t-1},z_t}}{c_{z_{t-1}}}$$

Transition Matrix

Definition

- These probabilities can be stored in a matrix called transition matrix of a Markov Chain. The number on row j column j' is the estimated probability $\hat{\mathbb{P}}\{j'|j\}$. If there are 3 tokens $\{1, 2, 3\}$, the transition matrix is the following.

$$\begin{bmatrix} \hat{\mathbb{P}}\{1|1\} & \hat{\mathbb{P}}\{2|1\} & \hat{\mathbb{P}}\{3|1\} \\ \hat{\mathbb{P}}\{1|2\} & \hat{\mathbb{P}}\{2|2\} & \hat{\mathbb{P}}\{3|2\} \\ \hat{\mathbb{P}}\{1|3\} & \hat{\mathbb{P}}\{2|3\} & \hat{\mathbb{P}}\{3|3\} \end{bmatrix}$$

- Given the initial distribution of tokens, the distribution of the next token can be found by multiplying it by the transition probabilities.

Aside: Stationary Probability

Discussion

- Given the bigram model, the fraction of times a token occurs for a document with infinite length can be computed. The resulting distribution is called the stationary distribution.

$$p_{\infty} = p_0 M^{\infty}$$

Aside: Spectral Decomposition

Discussion

- It is easier to find powers of diagonal matrices.
- Let D be the diagonal matrix with eigenvalues of M on the diagonal and P be the matrix with columns being corresponding eigenvectors.

$$MP = \lambda_i P, i = 1, 2, \dots, K$$

$$MP = PD$$

$$M = PDP^{-1}$$

$$M^n = \underbrace{PDP^{-1}PDP^{-1} \dots PDP^{-1}}_{n \text{ times}} = PD^n P^{-1}$$

$$M^\infty = PD^\infty P^{-1}$$

Aside: Stationarity

Discussion

- A simpler way to compute the stationary distribution is to solve the equation:

$$p_{\infty} = p_{\infty} M$$

Trigram Model

Definition

- The same formula can be applied to trigram: sequences of three tokens.

$$\hat{\mathbb{P}} \{z_t | z_{t-1}, z_{t-2}\} = \frac{c_{z_{t-2}, z_{t-1}, z_t}}{c_{z_{t-2}, z_{t-1}}}$$

- In a document, likely, these longer sequences of tokens never appear. In those cases, the probabilities are $\frac{0}{0}$. Because of this, Laplace smoothing adds 1 to all counts.

$$\hat{\mathbb{P}} \{z_t | z_{t-1}, z_{t-2}\} = \frac{c_{z_{t-2}, z_{t-1}, z_t} + 1}{c_{z_{t-2}, z_{t-1}} + m}$$

Laplace Smoothing

Definition

- Laplace smoothing should be used for bigram and unigram models too.

$$\hat{\mathbb{P}} \{z_t | z_{t-1}\} = \frac{c_{z_{t-1}, z_t} + 1}{c_{z_{t-1}} + m}$$

$$\hat{\mathbb{P}} \{z_t\} = \frac{c_{z_t} + 1}{\sum_{z=1}^m c_z + m}$$

- Aside: Laplace smoothing can also be used in decision tree training to compute entropy.

N Gram Model

Algorithm

- Input: series $\{z_1, z_2, \dots, z_{d_i}\}_{i=1}^n$.
- Output: transition probabilities $\hat{\mathbb{P}}\{z_t | z_{t-1}, z_{t-2}, \dots, z_{t-N+1}\}$ for all $z_t = 1, 2, \dots, m$.
- Compute the transition probabilities using counts and Laplace smoothing.

$$\hat{\mathbb{P}}\{z_t | z_{t-1}, z_{t-2}, \dots, z_{t-N+1}\} = \frac{c_{z_{t-N+1}, z_{t-N+2}, \dots, z_t} + 1}{c_{z_{t-N+1}, z_{t-N+2}, \dots, z_{t-1}} + m}$$

Sampling from Discrete Distribution

Discussion

- To generate new sentences given an N gram model, random realizations need to be generated given the conditional probability distribution.
- Given the first $N - 1$ words, z_1, z_2, \dots, z_{N-1} , the distribution of next word is approximated by $p_x = \hat{\mathbb{P}} \{z_N = x | z_{N-1}, z_{N-2}, \dots, z_1\}$. This process then can be repeated for on $z_2, z_3, \dots, z_{N-1}, z_N$ and so on.

Inverse Transform Sampling, Part I

Discussion

- Most programming languages have a function to generate a random number $u \sim \text{Unif}[0, 1]$.
- If there are $m = 2$ tokens in total and the conditional probabilities are p and $1 - p$. Then the following distributions are the same.

$$z_N = \begin{cases} 0 & \text{with probability } p \\ 1 & \text{with probability } 1 - p \end{cases} \Leftrightarrow z_N = \begin{cases} 0 & \text{if } 0 \leq u \leq p \\ 1 & \text{if } p < u \leq 1 \end{cases}$$

Inverse Transform Sampling, Part II

Discussion

- In the general case with m tokens with conditional probabilities p_1, p_2, \dots, p_m with $\sum_{j=1}^m p_j = 1$. Then the following distributions are the same.

$$z_N = j \text{ with probability } p_j \Leftrightarrow z_N = j \text{ if } \sum_{j'=1}^{j-1} p_{j'} < u \leq \sum_{j'=1}^j p_{j'}$$

- This can be used to generate a random token from the conditional distribution.

Sparse Matrix

Discussion

- The transition matrix is too large with mostly zeros.
- Usually, clustering is done so each type (or feature) represent a group of words.
- For the homework, treat each character (letter or space) as a token, then there are $26 + 1$ types. All punctuations are removed or converted to spaces.