
CS 232: Artificial Intelligence

Fall 2026

Prof. Carolyn Anderson
Wellesley College

Reminders

ONLINE


- ♦ I have help hours today (~~W423~~) and Monday from 3-4 (W323)
- ♦ Midterm 1 on Tuesday! I have notecards for you.
- ♦ Tutor hours:
 - Aeka: Thursdays 7:30-8:30 in H103
 - Amber: Friday 1:30-2:30 in Whiteboard Alley
 - Coco: Mondays 7-8 in Whiteboard Alley

Recap

Reinforcement Learning

Basic idea:

- ♦ Agent receives feedback in the form of rewards
- ♦ Agent tries to learn to act in order to maximize expected rewards
- ♦ Agent learns from samples of observed outcomes (AKA: try something out and see what happens!)



RL Components

- ✦ Policy selection: how do we pick actions?
- ✦ Policy evaluation: how do we learn from the results of our actions?

Policy Selection

RL policy selection must balance two things:

- ♦ **Exploitation**: choose the action with highest reward (according to current information)
- ♦ **Exploration**: choose a new action to discover new information more quickly.
- ♦ Different strategies for exploring:
 - Choose the action that has been tried least often in the current state.
 - Choose an action that leads to states that have been unexplored.
 - Choose a random action.

Semester Road Map

Almost all AI tasks can be grouped into one of three main categories:

- ♦ Search
- ♦ Classification
- ♦ Generation

AI Needs You Discussion

- ◆ What lessons does this chapter have for how we think about AI?
- ◆ What differences do you see between this era's technology debate and the current debate over AI?


Classification Methods

Classification Methods: Supervised Machine Learning

Lots of kinds!

- Naïve Bayes
- Logistic regression
- Neural networks
- k-Nearest Neighbors
- random forests
- ...

we'll cover these



Classification Methods: Supervised Machine Learning

Input:

- a input d
- a fixed set of classes $C = \{c_1, c_2, \dots, c_n\}$
- A training set of m hand-labeled examples

$$(d_1, c_1), \dots, (d_m, c_m)$$

Output:

- a learned classifier

$$f(d) \rightarrow c$$

Defining a Machine Learning Classifier

Given m input/output pairs $(x^{(i)}, y^{(i)})$:

1. A **feature representation** of the input
$$d = [x_1, x_2 \dots x_n]$$

x is our feature representation
2. A **classification function**
$$f(x) \rightarrow \hat{y} \quad p(y|x)$$
3. An objective function for learning, like **cross-entropy loss**
4. An algorithm for optimizing the objective function:
stochastic gradient descent.

Logistic Regression Classifiers

Is this spam?

Subject: Important notice!
From: Stanford University <newsforum@stanford.edu>
Date: October 28, 2011 12:34:16 PM PDT
To: undisclosed-recipients;;

Greats News!

You can now access the latest news by using the link below to login to Stanford University News Forum.

<http://www.123contactform.com/contact-form-StanfordNew1-236335.html>

Click on the above link to login for more information about this new exciting forum. You can also copy the above link to your browser bar and login for more information about the new services.

© Stanford University. All Rights Reserved.

Who wrote which Federalist papers?

Anonymous essays try to convince New York to ratify U.S Constitution.
Authorship of 12 of the letters in dispute.

Solved by Mosteller and Wallace (1963) using Bayesian methods



What is the subject of this research article?

MEDLINE Article



?

Blood Supply
Chemistry
Drug Therapy
Embryology
Epidemiology
...

Text Classification: definition

Input:

- a document d
- a fixed set of classes $C = \{c_1, \dots, c_n\}$
 $\{pos, neg\}$

Output: a predicted class

$$\hat{y} \in C$$

Sentiment example: does $y=1$ or $y=0$?

Leslie L.F 6 years ago

This is my favorite sandwich on earth. I usually use basil instead of chives/tarragon. Sometimes i add a sprinkle of shredded parmesean. So good. I love it so much.

[Reply](#) [Is this helpful?](#)  1

Kayla 7 years ago

This sandwich is outrageous. Make it now. I did it once begrudgingly and ate it 4 days that week for lunch.

[Is this helpful?](#)  7

Lynne 10 years ago

Really wanted to love this sandwich. Of course, I love the butter and the baguette (and the salt), but the radishes just don't do anything for me.

[Reply](#) [Is this helpful?](#) 

SonomaSMB 10 years ago

No! Butter is the essence of this. The radishes DO mind.

[Is this helpful?](#)  51

Radish Sandwiches With Butter and Salt

Recipe from "Root to Leaf"
Adapted by Sam Sifton
Updated June 10, 2021



Total Time 10 minutes
Rating 5 ★★★★★ (1,212)
Comments [Read 123 comments](#)

Steven Satterfield, the chef at Miller Union in Atlanta, included this very French picnic recipe in his cookbook, "Root to Leaf." As he points out, the key is to use a lot of butter, a lot of radishes and plenty of salt. The recipe yields four sturdy desk- or school-lunch sandwiches, or you can divide them further, into a dozen little bites for hors d'oeuvres. — Sam Sifton

franw 9 years ago

What a simple taste sensation! This is really a very filling and go-to sandwich when you are short on supply. That said, it makes for very pretty sandwiches when serving tea as well. Replace butter with cream cheese-wonderful!

The baguette is a very nice presentation but I have used Hawaiian sweet bread to shepherders and its all good. The versatility is key here.

I would recommend to try and save your judgement on this sandwich till you do.

[Is this helpful?](#)  12

 Save

Features in logistic regression

For feature x_i , weight w_i tells is how important is x_i

- x_i = "review contains **love**": $w_i = +10$
- x_j = "review contains **disappointing**": $w_j = -5$
- x_k = "review contains **outrageous**": $w_k = -2$

Logistic Regression for one observation x

Input observation: $x = [x_1 \dots x_n]$

Weights (one per feature): $W = [w_1 \dots w_n]$

Output: a predicted class $\hat{y} \in \{0, 1\}$

How to do classification

For each feature x_i , weight w_i tells us importance of x_i

bias : b

We'll sum up all the weighted features and the bias:

$$z = \left(\sum_{i=1}^n w_i x_i \right) + b$$

$$z = w \cdot x + b$$

But we want a probabilistic classifier

We need to formalize “sum is high”.

We want a model that can tell us:

$$p(y=1 \mid x; w)$$

$$p(y=0 \mid x; w)$$


The problem: z isn't a probability, it's just a number!

$$z = w \cdot x + b$$

$$\sigma(z) = \frac{1}{1 + e^{-z}}$$

sigmoid: outputs between 0 & 1

The very useful sigmoid or logistic function



0.5 = decision boundary

Idea of logistic regression

Compute $\frac{w \cdot x + b}{z}$

Then pass it through the sigmoid function:

$\sigma(w \cdot x + b)$

Treat it as a probability

Making probabilities with sigmoids

$$\begin{aligned} P(y = 1) &= \sigma(w \cdot x + b) \\ &= \frac{1}{1 + \exp(-(w \cdot x + b))} \end{aligned}$$

$$\begin{aligned} P(y = 0) &= 1 - \sigma(w \cdot x + b) \\ &= 1 - \frac{1}{1 + \exp(-(w \cdot x + b))} \end{aligned}$$


Turning a probability into a classifier

$$\text{decision}(x) = \begin{cases} 1 & \text{if } P(y = 1|x) > 0.5 \\ 0 & \text{otherwise} \end{cases}$$

0.5 is the decision boundary

Turning a probability into a classifier

0.5 is the **decision boundary**



The bias term b can be used to shift the intercept!

Logistic Regression Example: Text Classification

Sentiment example: does $y=1$ or $y=0$?

Leslie L.F 6 years ago

This is my favorite sandwich on earth. I usually use basil instead of chives/tarragon. Sometimes i add a sprinkle of shredded parmesean. So good. I love it so much.

[Reply](#) [Is this helpful?](#)  1

Kayla 7 years ago

This sandwich is outrageous. Make it now. I did it once begrudgingly and ate it 4 days that week for lunch.

[Is this helpful?](#)  7

Lynne 10 years ago

Really wanted to love this sandwich. Of course, I love the butter and the baguette (and the salt), but the radishes just don't do anything for me.

[Reply](#) [Is this helpful?](#) 

SonomaSMB 10 years ago

No! Butter is the essence of this. The radishes DO mind.

[Is this helpful?](#)  51

Radish Sandwiches With Butter and Salt

Recipe from "Root to Leaf"
Adapted by Sam Sifton
Updated June 10, 2021



Total Time 10 minutes
Rating 5 ★★★★★ (1,212)
Comments [Read 123 comments](#)

Steven Satterfield, the chef at Miller Union in Atlanta, included this very French picnic recipe in his cookbook, "Root to Leaf." As he points out, the key is to use a lot of butter, a lot of radishes and plenty of salt. The recipe yields four sturdy desk- or school-lunch sandwiches, or you can divide them further, into a dozen little bites for hors d'oeuvres. — Sam Sifton

franw 9 years ago

What a simple taste sensation! This is really a very filling and go-to sandwich when you are short on supply. That said, it makes for very pretty sandwiches when serving tea as well. Replace butter with cream cheese-wonderful!

The baguette is a very nice presentation but I have used Hawaiian sweet bread to shepherders and its all good. The versatility is key here.

I would recommend to try and save your judgement on this sandwich till you do.

[Is this helpful?](#)  12

 Save

Features

Var	Definition	Value
x_1	count(positive words)	2
x_2	count(negative words)	0
x_3	1 if "no" or "n't" in doc else 0	1
x_4	1 if ! in doc else 0	0
x_5	log(word count)	3.33

Weights

Var	Definition	Value	<u>W</u>
w_1	weight on positive words	2	2.5
w_2	weight on negative words	0	-2.5
w_3	weight on negation value	1	-3
w_4	weight on ! value	0	4
w_5	weight on length	3.33	0.1

$$b = 0.2$$

Classifying sentiment for input x

$$\begin{aligned} p(+1|x) &= \sigma(w \cdot x + b) \\ &= \sigma([2.5, -2.5, -3, 4, 0.1]^* \\ &\quad [2, 0, 1, 0, 3.33] + 0.2) \\ &= \sigma(2.53) \\ &= 0.926 \end{aligned}$$

$$p(-1|x) = 1 - 0.926 = 0.074$$

Classification in (binary) logistic regression: summary

Given:

- a set of classes: (+ sentiment, - sentiment)
- a vector $\mathbf{x}$ of features $[x_1, x_2, x_3, \dots, x_n]$
 - $x_1 = \text{count}(\text{"awesome"})$ in document
 - $x_3 = 1$ if "no" in document else 0
- A vector $\mathbf{w}$ of weights $[w_1, w_2, \dots, w_n]$
 - w_i for each feature f_i

$$\begin{aligned} P(y = 1) &= \sigma(w \cdot x + b) \\ &= \frac{1}{1 + \exp(-(w \cdot x + b))} \end{aligned}$$

Turning a probability into a classifier

$$\text{decision}(x) = \begin{cases} 1 & \text{if } P(y = 1|x) > 0.5 \\ 0 & \text{otherwise} \end{cases} \quad \begin{array}{l} \text{if } w \cdot x + b > 0 \\ \text{if } w \cdot x + b \leq 0 \end{array}$$

What do we conclude about this review?

$$p(y=1) = 92.6\%$$

Really wanted to love this sandwich. Of course, I love the butter and the baguette (and the salt), but the radishes just don't do anything for me.

Radish Sandwiches With Butter and Salt

Recipe from "Root to Leaf"
Adapted by Sam Sifton
Updated June 10, 2021

Total Time 10 minutes
Rating 5 ★★★★★ (1,212)
Comments Read 123 comments ↕



Steven Satterfield, the chef at Miller Union in Atlanta, included this very French picnic recipe in his cookbook, "Root to Leaf." As he points out, the key is to use a lot of butter, a lot of radishes and plenty of salt. The recipe yields four sturdy desk- or school-lunch sandwiches, or you can divide them further, into a dozen little bites for hors d'oeuvres. — Sam Sifton

The two phases of logistic regression

Training: we learn weights W & bias b
using stochastic gradient descent & our
Test: loss function.

Given a test example x & a set of
learned weights W, b , compute
 $p(y|x)$ & return the label w/ highest
probability.