MACHINE LEARNING BASICS

David Kauchak
CS159 Spring 2023

1

Admin

Assignment 6

pre-pre enrollment

No office hours Friday

2

Machine Learning is…

Machine learning, a branch of artificial intelligence, concerns the construction and study of systems that can learn from data.

WIKIPEDIA
The Free Encyclopedia

3

Machine Learning is…

Machine learning is programming computers to optimize a performance criterion using example data or past experience.

-- Ethem Alpaydin

The goal of machine learning is to develop methods that can automatically detect patterns in data, and then to use the uncovered patterns to predict future data or other outcomes of interest.

-- Kevin P. Murphy

The field of pattern recognition is concerned with the automatic discovery of regularities in data through the use of computer algorithms and with the use of these regularities to take actions.

-- Christopher M. Bishop

4

Machine Learning is…

Machine learning is about predicting the future based on the past.
-- Hal Daume III

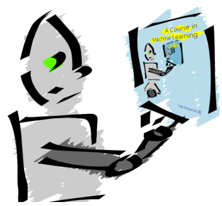

5

Machine Learning is…

Machine learning is about predicting the future based on the past.
-- Hal Daume III

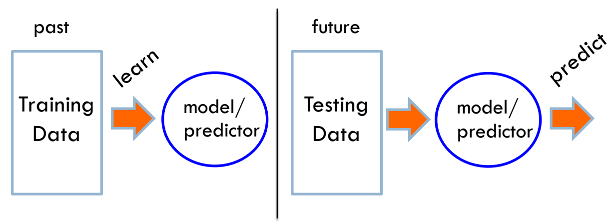

6

Why machine learning?

Lot's of data

Hand-written rules just don't do it

Performance can be much better than what people can do

Why not just study machine learning?

- Domain knowledge/expertise is still very important
- What types of features to use
- What models are important

7

Machine learning problems

What high-level machine learning problems and algorithms have you seen or heard of before?

8

Data

examples

Data

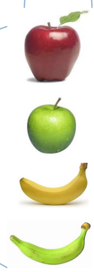

9

Data

examples

Data

10

Data

examples

Data

11

Data

examples

Data

12

Supervised learning

examples

label

label$_1$

label$_3$

label$_4$

label$_5$

labeled examples

Supervised learning: given labeled examples

13

Supervised learning

label

label$_1$

label$_3$

label$_4$

label$_5$

model/ predictor

Supervised learning: given labeled examples

14

Supervised learning

model/ predictor

predicted label

Supervised learning: learn to predict new example

15

Supervised learning: classification

label

apple

apple

banana

banana

Classification: a finite set of labels

Supervised learning: given labeled examples

16

NLP classification applications

Document classification
- spam
- sentiment analysis
- topic classification

Turn SafeSearch on or off
https://support.google.com/websearch/answer/510
1. Visit the Search Settings page.
2. In the "SafeSearch filters" section, select or unselect **Filter explicit results**.
3. Click **Save** at the bottom of the page.

Does linguistics phenomena X occur in text Y?

Digit recognition

Grammatically correct or not?

Word sense disambiguation

Any question you can pose as to have a discrete set of labels/answers!

17

Supervised learning: regression

label


-4.5


10.1

Regression: label is real-valued


3.2


4.3

Supervised learning: given labeled examples

18

Regression Example

Price of a used car

x : car attributes
(e.g. mileage)

y : price

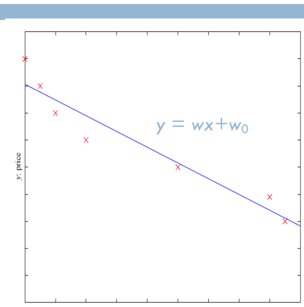

19

Regression applications

How many clicks will a particular website, ad, etc. get?

Predict the readability level of a document

Predict pause between spoken sentences?

Economics/Finance: predict the value of a stock

Car/plane navigation: angle of the steering wheel, acceleration, …

…

20

Supervised learning: ranking

label

1

4

2

3

Ranking: label is a ranking

Supervised learning: given labeled examples

21

NLP Ranking Applications

reranking N-best output lists (e.g. parsing, machine translation, …)

Rank possible simplification options

flight search (search in general)

…

22

Ranking example

Given a query and a set of web pages, rank them according to relevance

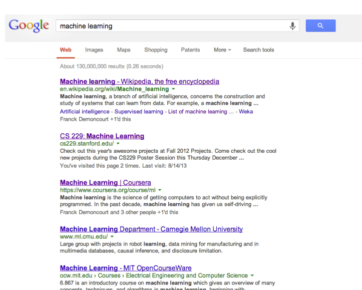

23

Unsupervised learning

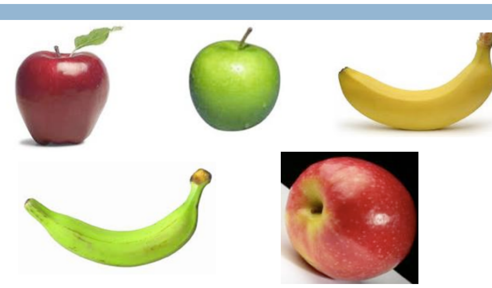

Unsupervised learning: given data, i.e. examples, but no labels

24

Unsupervised learning applications

learn clusters/groups without any label

- cluster documents
- cluster words (synonyms, parts of speech, …)

compression

bioinformatics: learn motifs

…

25

Reinforcement learning

left, right, straight, left, left, left, straight	GOOD
left, straight, straight, left, right, straight, straight	BAD

left, right, straight, left, left, left, straight	18.5
left, straight, straight, left, right, straight, straight	-3

Given a *sequence* of examples/states and a *reward* after completing that sequence, learn to predict the action to take in for an individual example/state

26

Reinforcement learning example

Backgammon

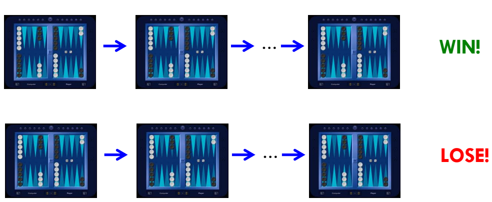

Given sequences of moves and whether or not the player won at the end, learn to make good moves

27

Reinforcement learning example

https://www.youtube.com/watch?v=tXlM99xPQC8

28

Other learning variations

What data is available:
- Supervised, unsupervised, reinforcement learning
- semi-supervised, active learning, …

How are we getting the data:
- online vs. offline learning

Type of model:
- generative vs. discriminative
- parametric vs. non-parametric

29

Text classification

label

spam

not spam

not spam

For this class, I'm mostly going to focus on classification

I'll use text classification as a running example

30

Representing examples

examples

What is an example?
How is it represented?

31

Features

examples

features

$f_1, f_2, f_3, \ldots, f_n$

$f_1, f_2, f_3, \ldots, f_n$

$f_1, f_2, f_3, \ldots, f_n$

$f_1, f_2, f_3, \ldots, f_n$

How our algorithms actually "view" the data

Features are the questions we can ask about the examples

32

Features

examples

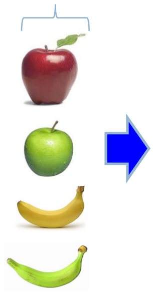

features

red, round, leaf, 3oz, …

green, round, no leaf, 4oz, …

yellow, curved, no leaf, 4oz, …

green, curved, no leaf, 5oz, …

How our algorithms actually “view” the data

Features are the questions we can ask about the examples

33

Text: raw data

Raw data

Features?

34

Feature examples

Raw data

Features

Clinton said banana repeatedly last week on tv, “banana, banana, banana”

(1, 1, 1, 0, 0, 1, 0, 0, …)

banana clinton said california across tv wrong capital

Occurrence of words (unigrams)

35

Feature examples

Raw data

Features

Clinton said banana repeatedly last week on tv, “banana, banana, banana”

(4, 1, 1, 0, 0, 1, 0, 0, …)

banana clinton said california across tv wrong capital

Frequency of word occurrence (unigram frequency)

36

Feature examples

Raw data

Features

Clinton said banana repeatedly last week on tv, "banana, banana, banana"

(1, 1, 1, 0, 0, 1, 0, 0, ...)

banana repeatedly
clinton said
said banana
california schools
across the
tv banana
wrong way
capital city

Occurrence of bigrams

37

Feature examples

Raw data

Features

Clinton said banana repeatedly last week on tv, "banana, banana, banana"

(1, 1, 1, 0, 0, 1, 0, 0, ...)

banana repeatedly
clinton said
said banana
california schools
across the
tv banana
wrong way
capital city

Other features?

38

Lots of other features

POS: occurrence, counts, sequence

Constituents

Whether 'V1agra' occurred 15 times

Whether 'banana' occurred more times than 'apple'

If the document has a number in it

…

Features are very important, but we're going to focus on the model

39

Classification revisited

examples	label
red, round, leaf, 3oz, …	apple
green, round, no leaf, 4oz, …	apple
yellow, curved, no leaf, 4oz, …	banana
green, curved, no leaf, 5oz, …	banana

learn

model/ classifier

During learning/training/induction, learn a model of what distinguishes apples and bananas *based on the features*

40

Classification revisited

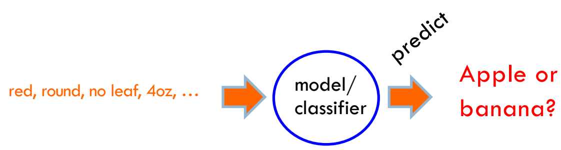

The model can then classify a new example *based on the features*

41

Classification revisited

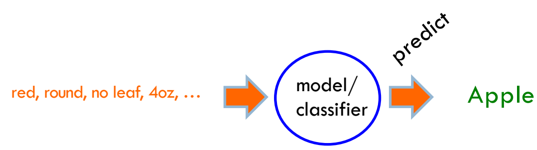

Why?

The model can then classify a new example *based on the features*

42

Classification revisited

Training data

examples	label
red, round, leaf, 3oz, …	apple
green, round, no leaf, 4oz, …	apple
yellow, curved, no leaf, 4oz, …	banana
green, curved, no leaf, 5oz, …	banana

Test set

red, round, no leaf, 4oz, … ?

43

Classification revisited

Training data

examples	label
red, round, leaf, 3oz, …	apple
green, round, no leaf, 4oz, …	apple
yellow, curved, no leaf, 4oz, …	banana
green, curved, no leaf, 5oz, …	banana

Test set

red, round, no leaf, 4oz, … ?

Learning is about ***generalizing*** from the training data

What does this assume about the training and test set?

44

Past predicts future

Training data

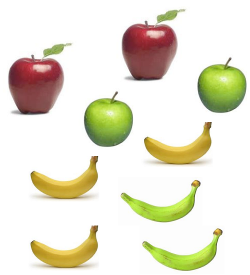

Test set

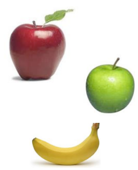

45

Past predicts future

Training data

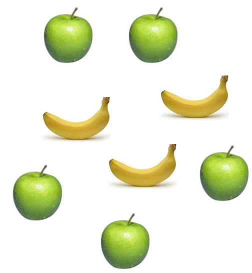

Test set

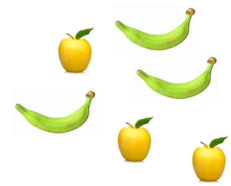

Not always the case, but we'll often assume it is!

46

Past predicts future

Training data

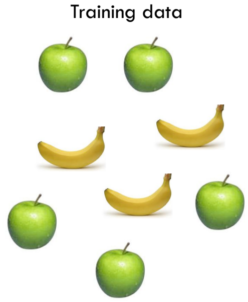

Test set

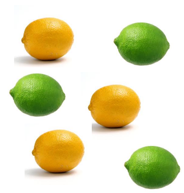

Not always the case, but we'll often assume it is!

47

More technically…

We are going to use the *probabilistic model* of learning

There is some probability distribution over example/label pairs called the *data generating distribution*

Both the training data **and** the test set are generated based on this distribution

48

data generating distribution

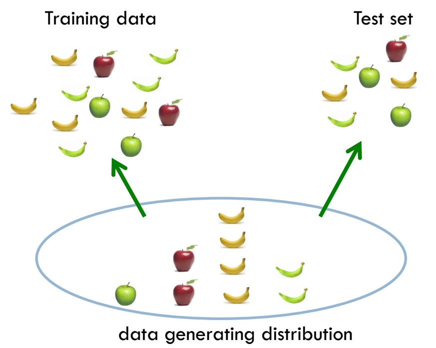

49

data generating distribution

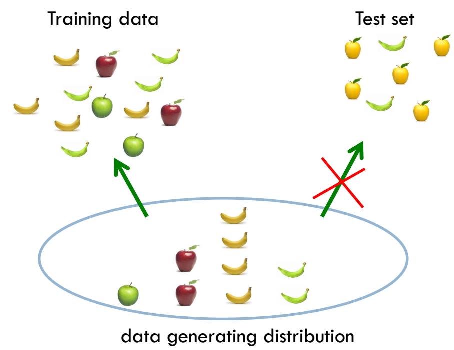

50

data generating distribution

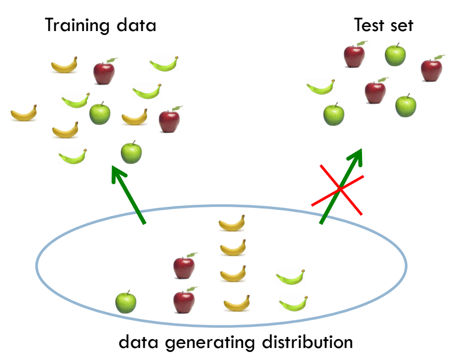

51

5 weeks left

Now that you know more about NLP, anything left that you'd like to know more about?

Summer plans?

Write the name of someone in the class that has their birthday later in the year than you?

52

Probabilistic Modeling

training data

train

probabilistic model

Model the data with a probabilistic model

specifically, learn *p(features, label)*

p(features, label) tells us how likely these features and this example are

53

An example: classifying fruit

Training data

examples	label
red, round, leaf, 3oz, …	apple
green, round, no leaf, 4oz, …	apple
yellow, curved, no leaf, 4oz, …	banana
green, curved, no leaf, 5oz, …	banana

train

probabilistic model:

p(*features, label*)

54

Probabilistic models

Probabilistic models define a *probability distribution* over features and labels:

yellow, curved, no leaf, 6oz, banana → probabilistic model: p(*features, label*) → 0.004

55

Probabilistic model vs. classifier

Probabilistic model:

yellow, curved, no leaf, 6oz, banana → probabilistic model: p(*features, label*) → 0.004

Classifier:

yellow, curved, no leaf, 6oz → probabilistic model: p(*features, label*) → banana

56

Probabilistic models: classification

Probabilistic models define a *probability distribution* over features and labels:

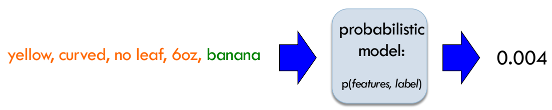

Given an unlabeled example: yellow, curved, no leaf, 6oz predict the label

How do we use a probabilistic model for classification/prediction?

57

Probabilistic models

Probabilistic models define a *probability distribution* over features and labels:

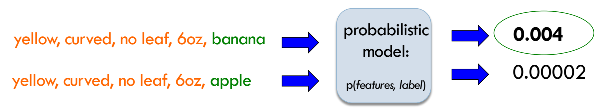

For each label, ask for the probability under the model
Pick the label with the highest probability

58

Probabilistic model vs. classifier

Probabilistic model:

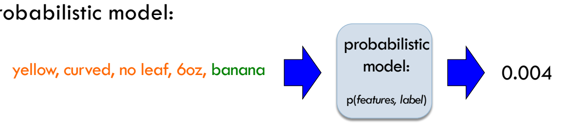

Classifier:

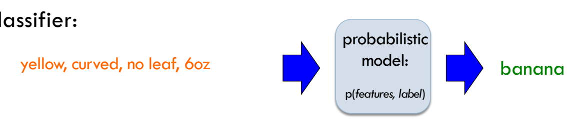

Why probabilistic models?

59

Probabilistic models

Probabilities are nice to work with
- range between 0 and 1
- can combine them in a well understood way
- lots of mathematical background/theory

Provide a strong, well-founded groundwork
- Allow us to make clear decisions about things like smoothing
- Tend to be much less "heuristic"
- Models have very clear meanings

60

Probabilistic models: big questions

1. Which model do we use, i.e. how do we calculate p(*feature, label*)?
2. How do train the model, i.e. how to we we estimate the probabilities for the model?
3. How do we deal with overfitting (i.e. smoothing)?

61

Basic steps for probabilistic modeling

Step 1: pick a model

Step 2: figure out how to estimate the probabilities for the model

Step 3 (optional): deal with overfitting

Probabilistic models

Which model do we use, i.e. how do we calculate p(*feature, label*)?

How do train the model, i.e. how to we we estimate the probabilities for the model?

How do we deal with overfitting?

62

What was the data generating distribution?

Training data

Test set

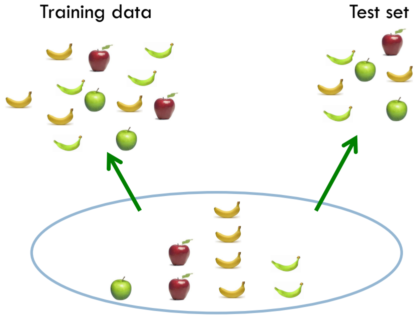

data generating distribution

63

Step 1: picking a model

What we're really trying to do is model the data generating distribution, that is how likely the feature/label combinations are

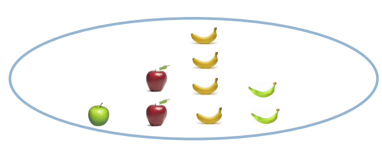

data generating distribution

64

Some math

$$p(features, label) = p(x_1, x_2, \ldots, x_m, y)$$

$$= p(y)p(x_1, x_2, \ldots, x_m | y)$$

What rule?

65

Some math

$$p(features, label) = p(x_1, x_2, \ldots, x_m, y)$$

$$= p(y)p(x_1, x_2, \ldots, x_m | y)$$

$$= p(y)p(x_1 | y)p(x_2, \ldots, x_m | y, x_1)$$

$$= p(y)p(x_1 | y)p(x_2 | y, x_1)p(x_3, \ldots, x_m | y, x_1, x_2)$$

$$= p(y)\prod_{j=1}^{m} p(x_i | y, x_1, \ldots, x_{i-1})$$

66

Step 1: pick a model

$$p(features, label) = p(y)\prod_{j=1}^{m} p(x_i | y, x_1, \ldots, x_{i-1})$$

So, far we have made NO assumptions about the data

$$p(x_m | y, x_1, x_2, \ldots, x_{m-1})$$

How many entries would the probability distribution table have if we tried to represent all possible values and we had 7000 binary features?

67

Full distribution tables

x_1	x_2	x_3	…	y	p( )
0	0	0	…	0	*
0	0	0	…	1	*
1	0	0	…	0	*
1	0	0	…	1	*
0	1	0	…	0	*
0	1	0	…	1	*
			…		

All possible combination of features!

Table size: 2^{7000} = ?

68

2^{7000}

1621696755662202026466665085478377095191112430363743256235982084151527023162702352987080237879
4460004651996019099530984538652557892546513204107022110253564658647431585227076599373340842842
7224200122818782600729310826170431944842663920777841250999968601694360066600112098175792966787
8196255237700655294757256678055809293844627218640216108862600816097132874749204352087401101862
6908423275017246052311293955235059054544214554772509509096507889478094683592939574112569473438
6191215296848474344406741204174020887540371869421701550220735398381224299258743537536161041593
4359455766656170179090417259702533652666268202180849389281269970952857089069637557541434487608
8248369941993802415197514510125127043829087280919538476302857811854024099958895964192277601255
3604911562403499947144160905730842429313962119953679373012944795600248333570738998392029910322
3465980389530690429801740098017325210691307971242016963397230218353007589784519525848553710885
8195631737000743805167411189134617501484521767984296782842287373127422122022517597535994839257
0298779077063553347902449354353866605125910795672914312162977887848185522928196541766009803989
9799168140474938421574351580260381151068286406789730483829220346042775765507377656754750702714
4662263487685709621261074762705203049488907208978593689047063428548531668665657327174660658185
6090664849508012761754614572161769555751992117507514067775104496728590822558547771447242334900
7640263217608921135525612411945387026802990440018385850576719369689759366121356888838680023840
9325673807775018914703049621509969838539752071549396339237202875920415172949370790977853625108
3200928396048072379548870695466216880446521124930762900919907177423550391351174415329737479300
8995583051888413533479846411368000499940373724560035428811232632821866113106455077289922996946
9156018580839820741704606832124388152026099584696588161375826382921029547343888832163627122302
9212297953848683554835357106034077891774170263636562027269554375177807413134551018100094688094
0781122057380335371124632958916237089580476224595091825301636909236240671411644331656159828058
372078343988856239089202844090255382937

Any problems with this?

69

Full distribution tables

x_1	x_2	x_3	…	y	p( )
0	0	0	…	0	*
0	0	0	…	1	*
1	0	0	…	0	*
1	0	0	…	1	*
0	1	0	…	0	*
0	1	0	…	1	*
			…		

- Storing a table of that size is impossible!
- How are we supposed to learn/estimate each entry in the table?

70

Step 1: pick a model

$$p(features, label) = p(y)\prod_{j=1}^{m} p(x_i \mid y, x_1, \ldots, x_{i-1})$$

So, far we have made NO assumptions about the data

Model selection involves making assumptions about the data

We've done this before, n-gram language model, parsing, etc.

These assumptions allow us to represent the data more compactly and to estimate the parameters of the model

71

Naïve Bayes assumption

$$p(features, label) = p(y)\prod_{j=1}^{m} p(x_i \mid y, x_1, \ldots, x_{i-1})$$

$$p(x_i \mid y, x_1, x_2, \ldots, x_{i-1}) = p(x_i \mid y)$$

What does this assume?

72

Naïve Bayes assumption

$$p(features, label) = p(y)\prod_{j=1}^{m} p(x_i \mid y, x_1, \ldots, x_{i-1})$$

$$p(x_i \mid y, x_1, x_2, \ldots, x_{i-1}) = p(x_i \mid y)$$

Assumes feature i is independent of the the other features *given the label*

Is this true for text, say, with unigram features?

73

Naïve Bayes assumption

$$p(x_i \mid y, x_1, x_2, \ldots, x_{i-1}) = p(x_i \mid y)$$

For most applications, this is not true!

For example, the fact that "San" occurs will probably make it *more likely* that "Francisco" occurs

However, this is often a reasonable approximation:

$$p(x_i \mid y, x_1, x_2, \ldots, x_{i-1}) \approx p(x_i \mid y)$$

74