GEOMETRIC VIEW OF DATA

David Kauchak
CS 158 – Fall 2025

1

Admin

- Assignment 2 out and due on Sunday
- Assignment 1 solution posted under the "File" tab on canvas (use them to debug!)
- Assignment 1 back soon
- Keep reading
- Mentor hours: Wednesday, 6-8pm (Alan, Edmunds downstairs)
- Office hours:
  - Mon, Tue, Thu: 10-11am
  - Thu: 4-5pm

2

Proper Experimentation

3

Experimental setup

REAL WORLD USE OF ML ALGORITHMS

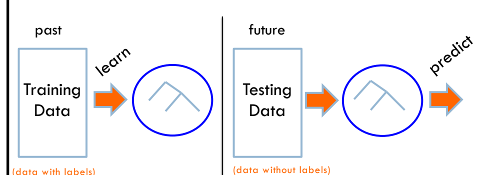

How do we tell how well we're doing?

4

Real-world classification

Google has labeled training data, for example from people clicking the "spam" button, but when new messages come in, they're not labeled

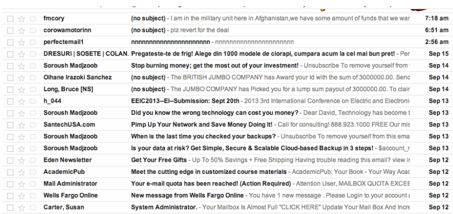

5

Classification evaluation

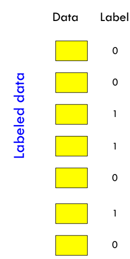

Use the labeled data we have already to create a test set with known labels!

Why can we do this?

We assume there's an underlying distribution that generates both the training and test examples

6

Classification evaluation

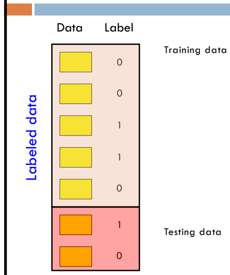

7

Classification evaluation

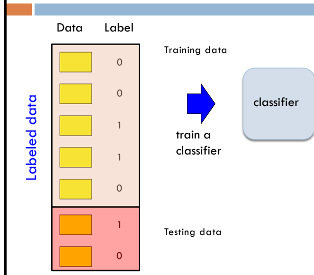

8

Classification evaluation

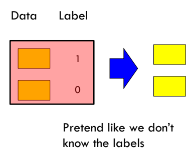

9

Classification evaluation

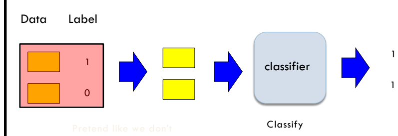

10

Classification evaluation

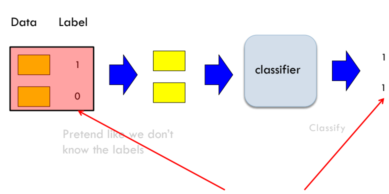

How could we score these for classification?

Compare predicted labels to actual labels

11

Test accuracy

To evaluate the model, compare the predicted labels to the actual labels

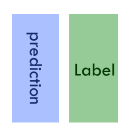

Accuracy: the proportion of examples where we correctly predicted the label

12

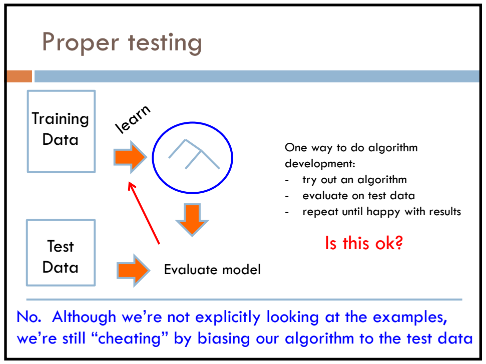
Proper testing
Training Data
learn
Test Data
Evaluate model
One way to do algorithm development:
- try out an algorithm
- evaluate on test data
- repeat until happy with results
Is this ok?
No. Although we're not explicitly looking at the examples, we're still "cheating" by biasing our algorithm to the test data

13

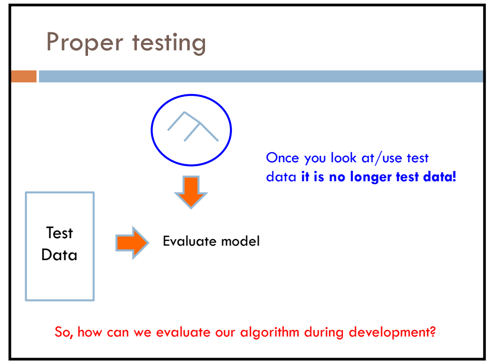
Proper testing
Once you look at/use test data it is no longer test data!
Test Data
Evaluate model
So, how can we evaluate our algorithm during development?

14

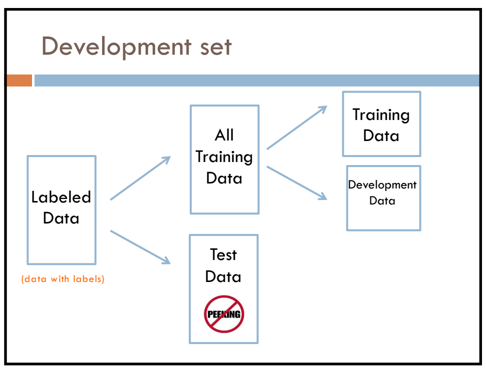
Development set
Labeled Data
(data with labels)
All Training Data
Training Data
Development Data
Test Data
PEEKING

15

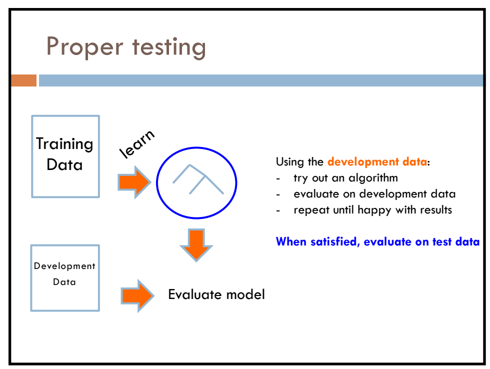
Proper testing
Training Data
learn
Development Data
Evaluate model
Using the development data:
- try out an algorithm
- evaluate on development data
- repeat until happy with results
When satisfied, evaluate on test data

16

Proper testing

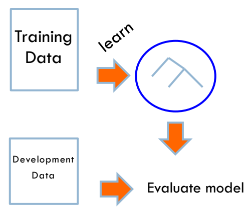

Using the **development data:**
- try out an algorithm
- evaluate on development data
- repeat until happy with results

Any problems with this?

17

Overfitting to development data

Be careful not to overfit to the development data!

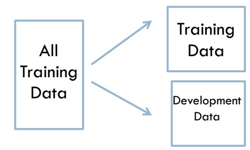

Often we'll split off development data multiple times (in fact, on the fly)… you can still overfit, but this helps avoid it

18

Pruning revisited

Which should we pick?

19

Pruning revisited

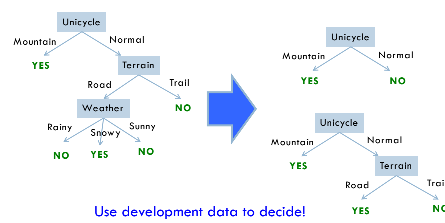

Use development data to decide!

20

Machine Learning: A Geometric View

21

Apples vs. Bananas

Weight	Color	Label
4	Red	Apple
5	Yellow	Apple
6	Yellow	Banana
3	Red	Apple
7	Yellow	Banana
8	Yellow	Banana
6	Yellow	Apple

Can we visualize this data?

22

Apples vs. Bananas

Turn features into numerical values
(read the book for a more detailed discussion of this)

Weight	Color	Label
4	0	Apple
5	1	Apple
6	1	Banana
3	0	Apple
7	1	Banana
8	1	Banana
6	1	Apple

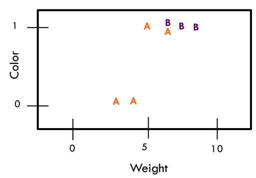

We can view examples as points in an *n*-dimensional space where *n* is the number of features

23

Examples in a feature space

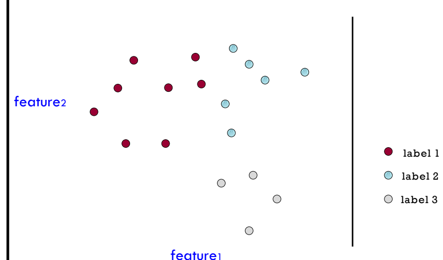

24

Test example: what class?

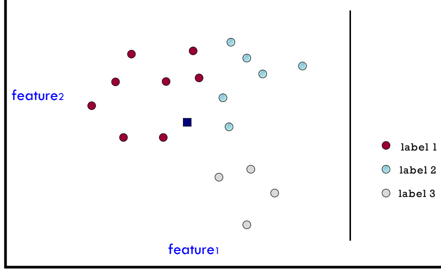

25

Test example: what class?

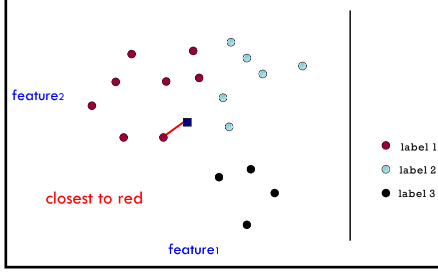

26

Another classification algorithm?

To classify an example ***d***:

Label ***d*** with the label of the closest example to ***d*** in the training set

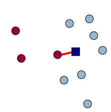

27

What about this example?

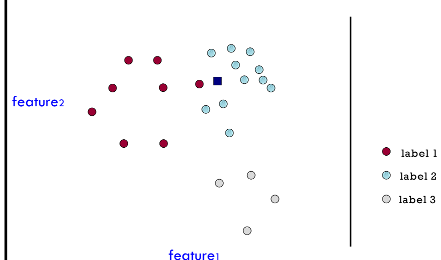

28

What about this example?

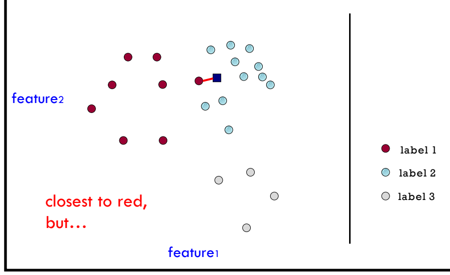

29

What about this example?

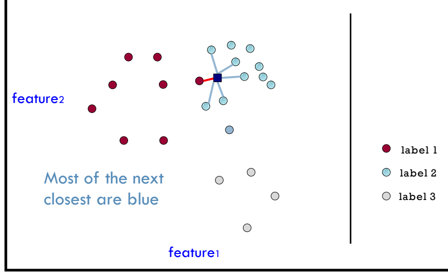

30

k-Nearest Neighbor (k-NN)

To classify an example ***d***:

- Find ***k*** nearest neighbors of ***d***
- Choose as the label the majority label within the ***k*** nearest neighbors

31

k-Nearest Neighbor (k-NN)

To classify an example ***d***:

- Find ***k*** *nearest* neighbors of ***d***
- Choose as the label the majority label within the ***k*** nearest neighbors

How do we measure "nearest"?

32

Euclidean distance

In two dimensions, how do we compute the distance?

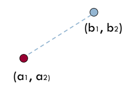

$$D(a,b)=\sqrt{(a_1-b_1)^2+(a_2-b_2)^2}$$

33

Euclidean distance

In n-dimensions, how do we compute the distance?

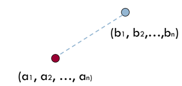

$$D(a,b)=\sqrt{(a_1-b_1)^2+(a_2-b_2)^2+\ldots+(a_n-b_n)^2}$$

34

Euclidean distance

In n-dimensions, how do we compute the distance?

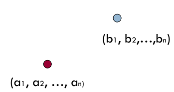

Measuring distance/similarity is a domain-specific problem and there are many, many different variations!

35

Decision boundaries

The **decision boundaries** are places in the features space where the classification of a point/example changes

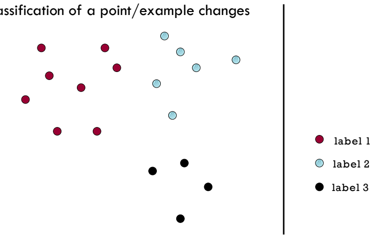

Where are the decision boundaries for k-NN?

36

k-NN decision boundaries

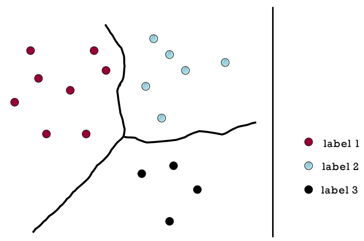

k-NN gives locally defined decision boundaries between classes

37

Choosing k

What is the label with k = 1?

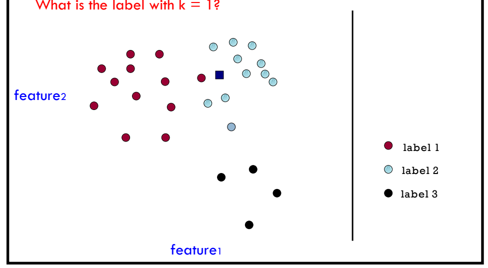

38

Choosing k

We'd choose red. Do you agree?

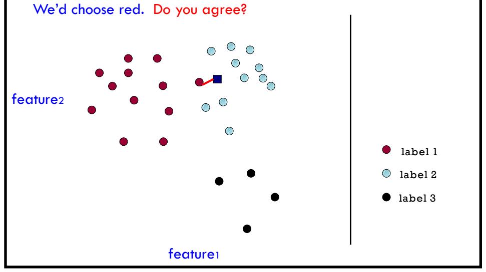

39

Choosing k

What is the label with k = 3?

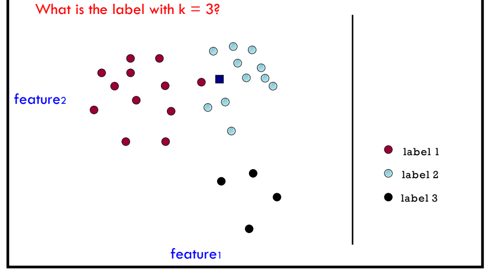

40

Choosing k

We'd choose blue. Do you agree?

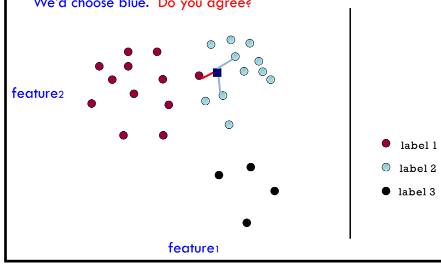

41

Choosing k

What is the label with k = 100?

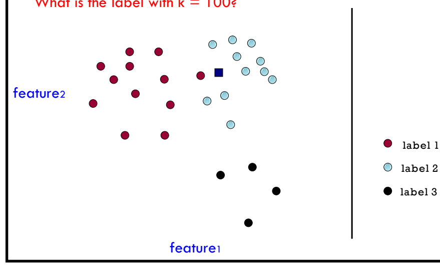

42

Choosing k

We'd choose red. Do you agree?

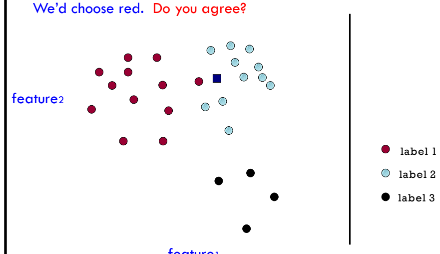

43

The impact of k

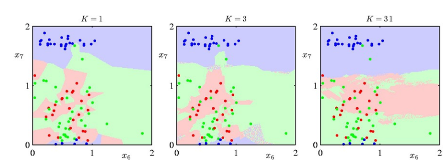

What is the role of k?
How does it relate to overfitting and underfitting?
How did we control this for decision trees?

44

k-Nearest Neighbor (k-NN)

To classify an example ***d***:

- Find ***k*** nearest neighbors of ***d***
- Choose as the class the majority class within the ***k*** nearest neighbors

How do we choose *k*?

45

How to pick k

Common heuristics:

- often 3, 5, 7
- choose an odd number to avoid ties

Use development data

46

k-NN variants

To classify an example ***d***:

- Find ***k*** nearest neighbors of ***d***
- Choose as the class the majority class within the ***k*** nearest neighbors

Any variation ideas?

47

k-NN variations

Instead of *k* nearest neighbors, count majority from all examples within a fixed distance

Weighted *k*-NN:

- Right now, all examples are treated equally
- weight the "vote" of the examples, so that closer examples have more vote/weight
- often use some sort of exponential decay

48

Decision boundaries for decision trees

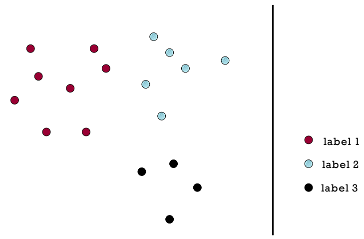

What do the decision boundaries for decision trees like?

49

Decision boundaries for decision trees

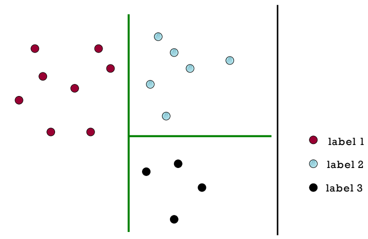

Axis-aligned splits/cuts of the data

50

Decision boundaries for decision trees

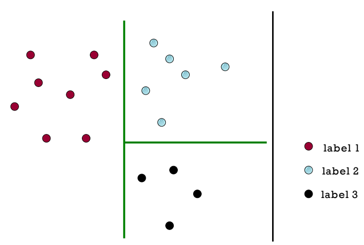

What types of data sets will DT work poorly on?

51

Problems for DT

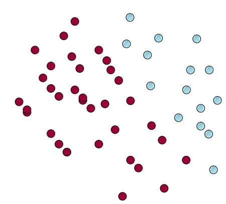

52

Decision trees vs. *k*-NN

Which is faster to train?

Which is faster to classify?

Do they use the features in the same way to label the examples?

53

Decision trees vs. *k*-NN

Which is faster to train?

k-NN doesn't require any training!

Which is faster to classify?

For most data sets, decision trees

Do they use the features in the same way to label the examples?

k-NN treats all features equally! Decision trees "select" important features

54

Machine learning models

Some machine learning approaches make strong assumptions about the data
- If the assumptions are true it can often lead to better performance
- If the assumptions aren't true, the approach can fail miserably

Other approaches don't make many assumptions about the data
- This can allow us to learn from more varied data
- But, they are more prone to overfitting
- and generally require more training data

55

Data generating distribution

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

What is a probability distribution?

56

Probability distribution

Describes how likely (i.e. probable) certain events are

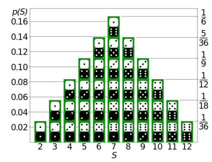

- Describes probabilities for all possible events
- Probabilities are between 0 and 1 (inclusive)
- Sum of probabilities over all events is 1

57

data generating distribution

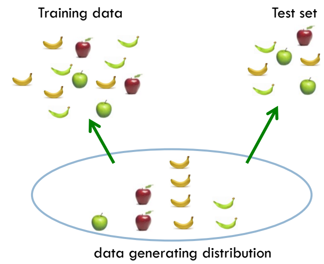

58

What is the data generating distribution?

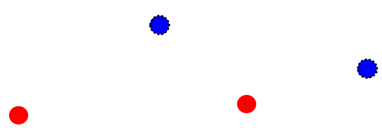

59

What is the data generating distribution?

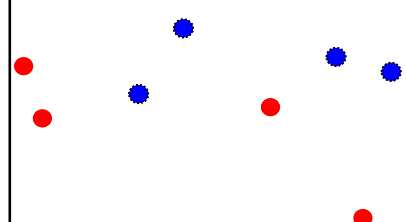

60

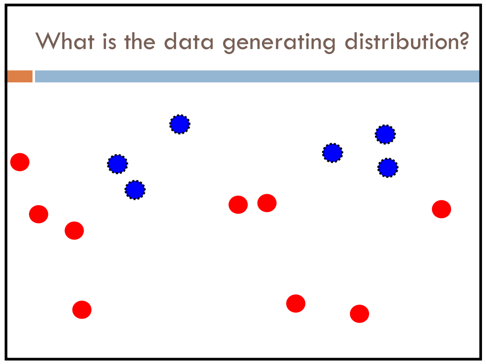
What is the data generating distribution?

61

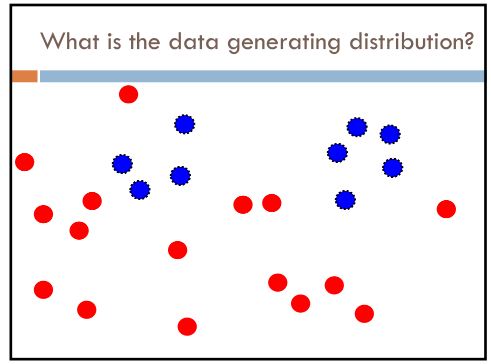
What is the data generating distribution?

62

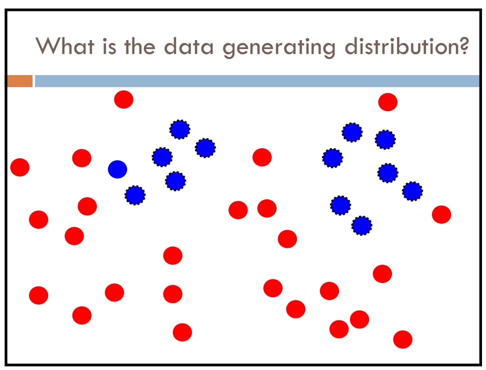
What is the data generating distribution?

63

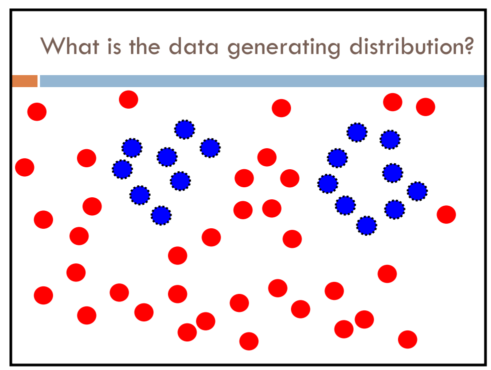
What is the data generating distribution?

64

Actual model

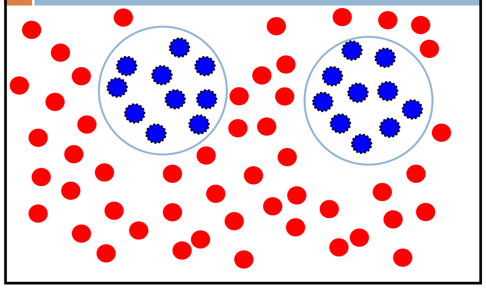

65

Model assumptions

If you don't have strong assumptions about the model, it can take you a longer to learn

Assume now that our model of the blue class is two circles

66

What is the data generating distribution?

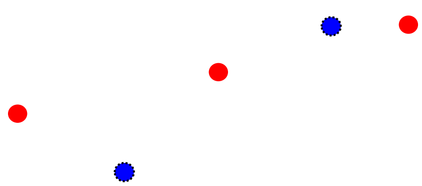

67

What is the data generating distribution?

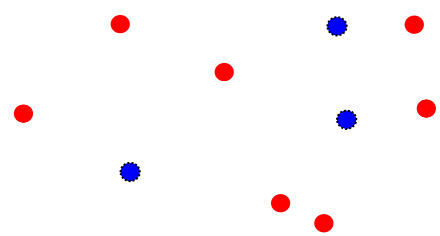

68

What is the data generating distribution?

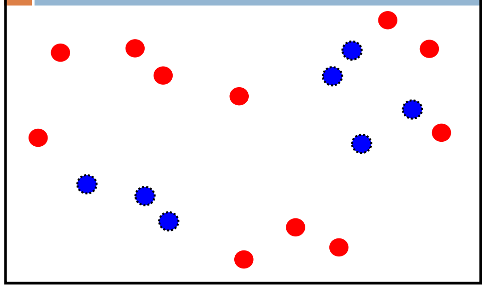

69

What is the data generating distribution?

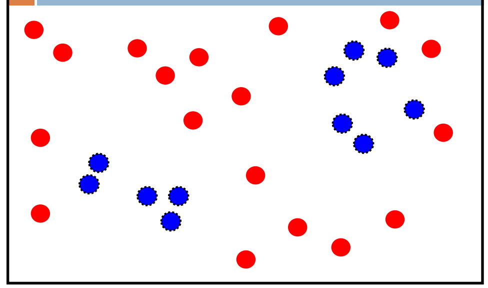

70

What is the data generating distribution?

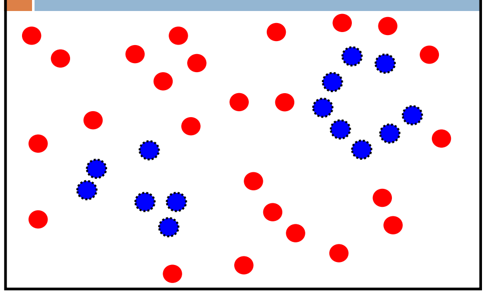

71

Actual model

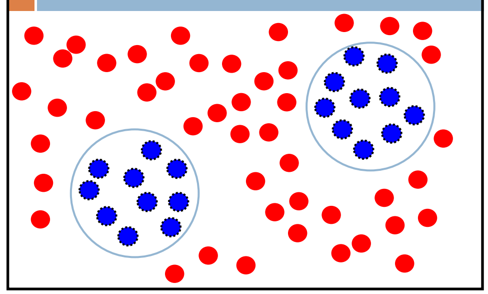

72

What is the data generating distribution?

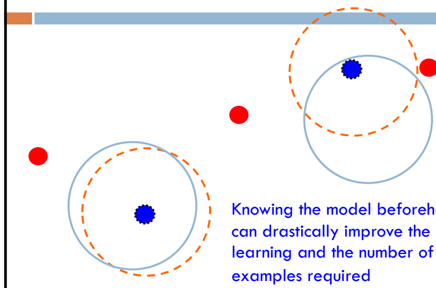

Knowing the model beforehand can drastically improve the learning and the number of examples required

73

What is the data generating distribution?

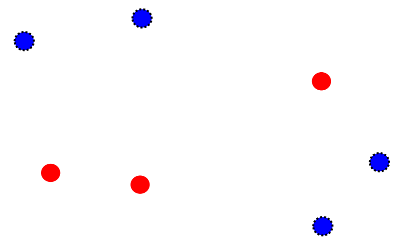

74

Make sure your assumption is correct, though!

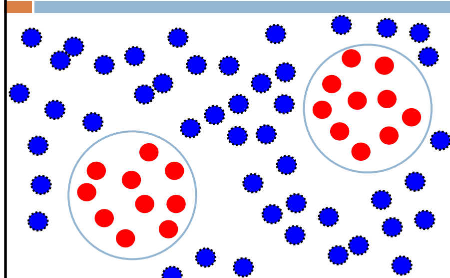

75

Machine learning models

What are the *model* assumptions (if any) that *k*-NN and decision trees make about the data?

Are there data sets that could never be learned correctly by either?

76

k-NN model

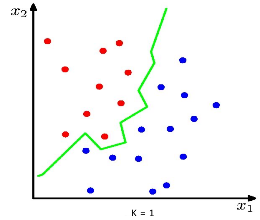

No model assumptions. Assumes that proximity relates to class

77

Decision tree model

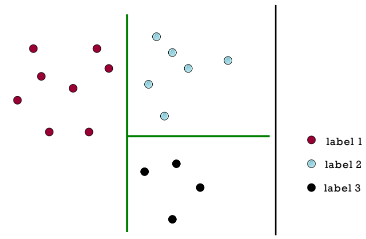

Axis-aligned splits/cuts of the data

78

Bias

The "bias" of a model is how strong the model assumptions are.

low-bias classifiers make minimal assumptions about the data (*k*-NN and DT are generally considered low bias)

high-bias classifiers make strong assumptions about the data

79

Linear models

A strong high-bias assumption is *linear separability*:
- in 2 dimensions, can separate classes by a line
- in higher dimensions, need hyperplanes

A *linear model* is a model that assumes the data is linearly separable

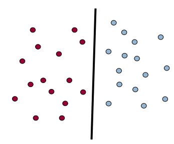

80

An aside: a thought experiment

What is a 100,000-dimensional space like?

You're a 1-D creature, and you decide to buy a 2-unit apartment

2 rooms (very, skinny rooms)

81

Another thought experiment

What is a 100,000-dimensional space like?

Your job's going well and you're making good money. You upgrade to a 2-D apartment with 2-units per dimension

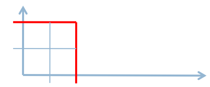

4 rooms (very, flat rooms)

82

Another thought experiment

What is a 100,000-dimensional space like?

You get promoted again and start having kids and decide to upgrade to another dimension.

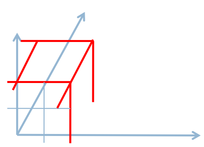

8 rooms (very, normal rooms)

Each time you add a dimension, the amount of space you have to work with goes up exponentially

83

Another thought experiment

What is a 100,000-dimensional space like?

Sundar Pichai steps down as CEO of Google (Alphabet) and they ask you if you'd like the job. You decide to upgrade to a 100,000 dimensional apartment.

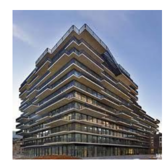

How much room do you have?
Can you have a big party?

$2^{100,000}$ rooms (it's very quiet and lonely…) = ~10^{30} rooms per person if you invited everyone on the planet

84

The challenge

Our intuitions about space/distance don't scale with dimensions!

85