

CS50

Machine Learning

Week 7

Machine Learning



what society thinks I
do




what my friends think
I do



what my parents think
I do

$$\begin{aligned} L_r &= \frac{1}{2} \|\mathbf{w}\|^2 - \sum_{i=1}^n \alpha_i y_i (\mathbf{x}_i \cdot \mathbf{w} + b) + \sum_{i=1}^n \alpha_i \\ \alpha_i &\geq 0, \forall i \\ \mathbf{w} &= \sum_{i=1}^n \alpha_i y_i \mathbf{x}_i, \sum_{i=1}^n \alpha_i y_i = 0 \\ \nabla \hat{g}(\theta_t) &= \frac{1}{n} \sum_{i=1}^n \nabla \ell(x_i, y_i; \theta_t) + \nabla r(\theta_t) \\ \theta_{t+1} &= \theta_t - \eta_t \nabla \ell(x_{i(t)}, y_{i(t)}; \theta_t) - \eta_t \cdot \nabla r(\theta_t) \\ \mathbb{E}_{i(t)}[\ell(x_{i(t)}, y_{i(t)}; \theta_t)] &= \frac{1}{n} \sum_i \ell(x_i, y_i; \theta_t). \end{aligned}$$

what other programmers
think I do

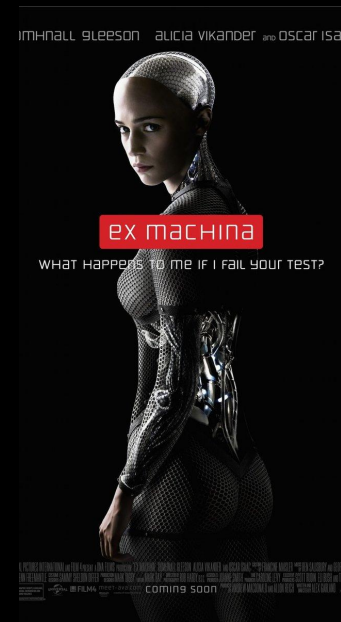



what I think I do


```
>>> from sklearn import svm
```

what I really do

Machine Learning?



Machine Learning?



Search Engines





Image Recognition






Voice Recognition



Natural Language Processing



Image Recognition




horse
car

Natural Language Processing

Nineteen Eighty-Four
by George Orwell
(1984)

[...]
BIG BROTHER
IS WATCHING
YOU, the caption
said, while the
dark eyes looked
deep into
Winston's own
[...]




Politics
Propaganda
Privacy

Whodunit!



Image recognition



Machine Learning algorithms



Machine Learning algorithms

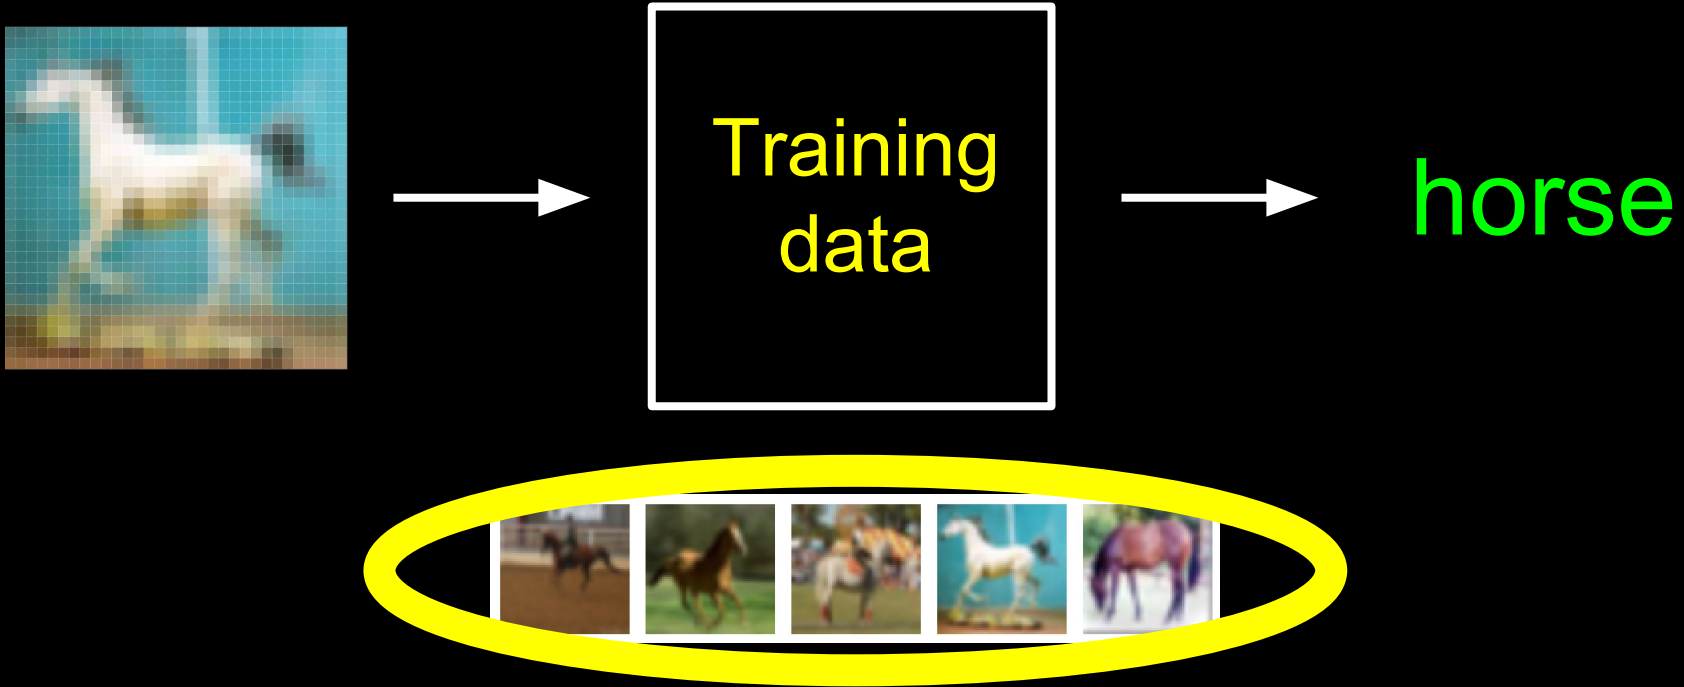




Image Classification

0
1
2
3
4
5
6
7
8
9



Handwritten digit classification




0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1
2	2	2	2	2	2	2	2	2	2	2	2	2	2	2	2	2	2	2	2
3	3	3	3	3	3	3	3	3	3	3	3	3	3	3	3	3	3	3	3
4	4	4	4	4	4	4	4	4	4	4	4	4	4	4	4	4	4	4	4
5	5	5	5	5	5	5	5	5	5	5	5	5	5	5	5	5	5	5	5
6	6	6	6	6	6	6	6	6	6	6	6	6	6	6	6	6	6	6	6
7	7	7	7	7	7	7	7	7	7	7	7	7	7	7	7	7	7	7	7
8	8	8	8	8	8	8	8	8	8	8	8	8	8	8	8	8	8	8	8
9	9	9	9	9	9	9	9	9	9	9	9	9	9	9	9	9	9	9	9

0
6

Nearest
Neighbor
Classifier

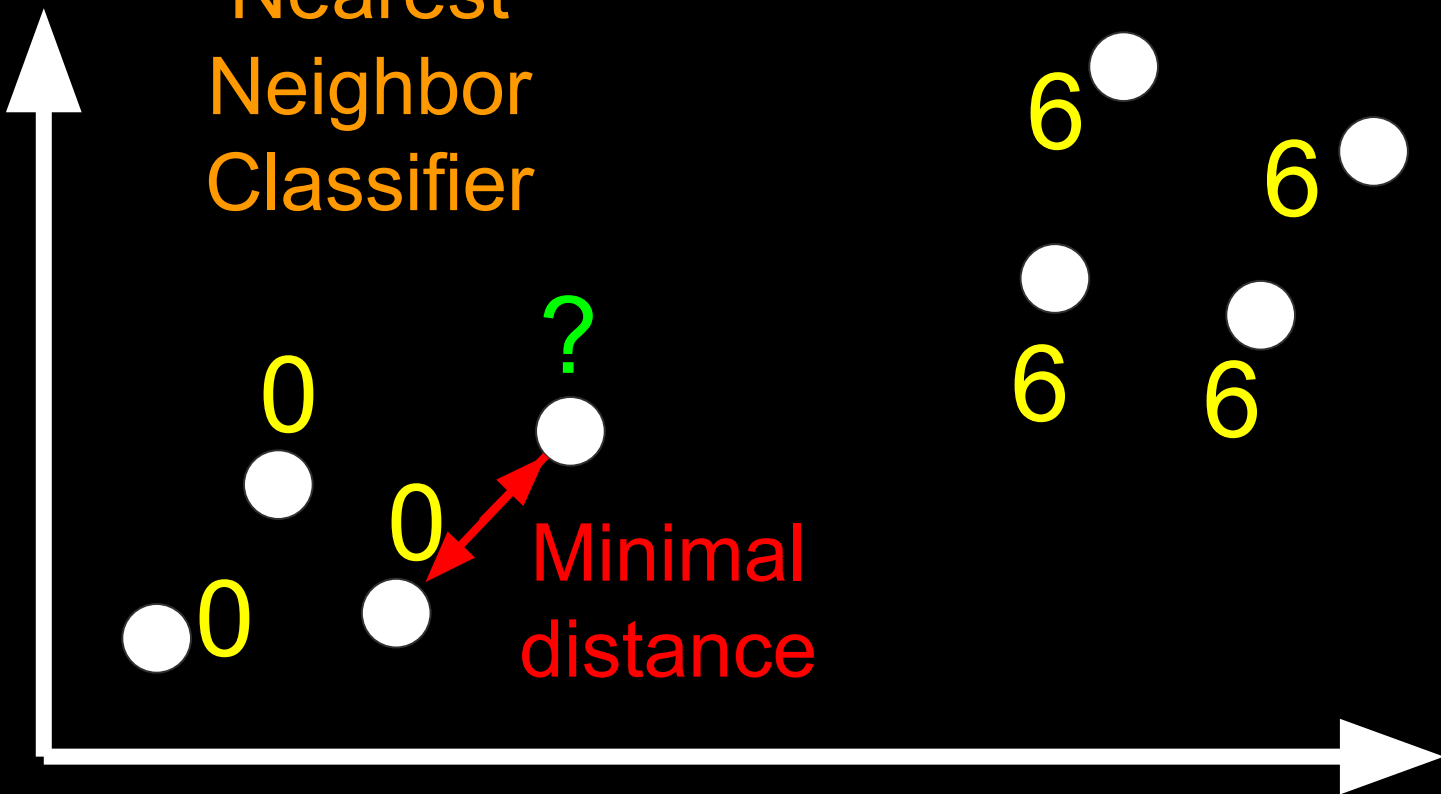
Minimal
distance



Labeled training set



Test point

Nearest Neighbor Classifier





Labeled training set

Test point





Nearest Neighbor Classifier







11000011
10111101
01011010
01111110
01011010
01100110
10111101
11000011



W.A.

"O day and night, but this is wondrous strange"



Flatland
by Edwin Abbott Abbott
(1884)



Flatland, Edwin Abbott Abbott, 1984

Flatland: The story describes a **two-dimensional world** occupied by geometric figures. The narrator is a square named **A Square** who guides the readers through some of the implications of life in two dimensions.



On New Year's Eve, A Square dreams about a visit to a **one-dimensional** world (**Lineland**) inhabited by "lustrous **points**", in which he attempts to convince the realm's monarch of a second dimension; but is unable to do so.

Following this vision, A Square is himself visited by a **three-dimensional** sphere named **A Sphere**, which he cannot comprehend until he sees **Spaceland** (a tridimensional world)


Ready to go beyond
Lineland, Flatland, and
Spaceland?




11000011
10111101
01011010
01111110
01011010
01100110
10111101
11000011




0	0	5	13	9	1	0	0
0	0	13	15	10	15	5	0
0	3	15	2	0	11	8	0
0	4	12	0	0	8	8	0
0	5	8	0	0	9	8	0
0	4	11	0	1	12	7	0
0	2	14	5	10	12	0	0
0	0	6	13	10	0	0	0






```
0 0 5 13 9 1 0 0
0 0 13 15 10 15 5 0
0 3 15 2 0 11 8 0
0 4 12 0 0 8 8 0
0 5 8 0 0 9 8 0
0 4 11 0 1 12 7 0
0 2 14 5 10 12 0 0
0 0 6 13 10 0 0 0
```



64 dimensional space

Nearest Neighbor Classifier





dist(


0	0	5	13	9	1	0	0
0	0	13	15	10	15	5	0
0	3	15	2	0	11	8	0
0	4	12	0	0	8	8	0
0	5	8	0	0	9	8	0
0	4	11	0	1	12	7	0
0	2	14	5	10	12	0	0
0	0	6	13	10	0	0	0

,

0	0	2	12	4	0	0	0
0	1	12	16	16	3	0	0
0	7	16	6	4	13	0	0
0	8	16	6	0	13	5	0
0	1	16	5	0	7	9	0
0	0	16	8	0	8	12	0
0	0	13	14	14	16	10	0
0	0	4	14	15	7	0	0

) = 31.98

The figure shows two 8x8 grayscale images of handwritten digits '0' and '1' with axes labeled 0-7. Below the images are two 8x8 distance matrices. The left matrix is a distance matrix for the '0' image, and the right matrix is a distance matrix for the '1' image. The matrices are separated by a red comma. A large red opening parenthesis is to the left of the first matrix, and a large red closing parenthesis is to the right of the second matrix. The result of the distance calculation is 31.98.



dist(

0	0	5	13	9	1	0	0	0	0	4	14	5	0	0	0
0	0	13	15	10	15	5	0	0	0	13	14	0	0	0	0
0	3	15	2	0	11	8	0	0	2	16	10	0	0	0	0
0	4	12	0	0	8	8	0	0	4	16	7	0	0	0	0
0	5	8	0	0	9	8	0	0	6	16	16	15	4	0	0
0	4	11	0	1	12	7	0	0	4	16	9	4	16	2	0
0	2	14	5	10	12	0	0	0	1	15	13	6	16	11	0
0	0	6	13	10	0	0	0	0	0	4	13	16	15	5	0

,

) = 45.97

The digits dataset

0
1
2
3
4
5
6
7
8
9



Labeled
training
set

Python code (Supervised Learning)

`np.sqrt(np.sum((x - y)**2))` ???

`x = [1, 1]`

`y = [3, 4]`

`x - y = [-2, -3]`

`(x - y)**2 = [4, 9]`

`np.sum((x - y)**2) = 13`

`np.sqrt(np.sum((x - y)**2)) = 3.60`



0
1
2
3
4
5
6
7
8
9



Labeled
training
set

Test point

Labeled
Training subset



0
1
2
3
4
5
6
7
8
9



Labeled
Training set


Testing set

0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1
2	2	2	2	2	2	2	2	2	2	2	2	2	2	2	2	2	2	2	2
3	3	3	3	3	3	3	3	3	3	3	3	3	3	3	3	3	3	3	3
4	4	4	4	4	4	4	4	4	4	4	4	4	4	4	4	4	4	4	4
5	5	5	5	5	5	5	5	5	5	5	5	5	5	5	5	5	5	5	5
6	6	6	6	6	6	6	6	6	6	6	6	6	6	6	6	6	6	6	6
7	7	7	7	7	7	7	7	7	7	7	7	7	7	7	7	7	7	7	7
8	8	8	8	8	8	8	8	8	8	8	8	8	8	8	8	8	8	8	8
9	9	9	9	9	9	9	9	9	9	9	9	9	9	9	9	9	9	9	9

Labeled
Training set

Testing set


With Nearest Neighbor Classifier



~ 97% Correct

The CIFAR-10 dataset


airplane
automobile
bird
cat
deer
dog
frog
horse
ship
truck



Labeled
training set


www.kaggle.com

With Nearest Neighbor Classifier




~ 30% Correct

Training set for category '0':



Training set for category 'horse':



Challenges

Viewpoint variation



Scale variation



Deformation



Occlusion



Illumination conditions



Background clutter




Intra-class variation





Features

0	0	5	13	9	1	0	0
0	0	13	15	10	15	5	0
0	3	15	2	0	11	8	0
0	4	12	0	0	8	8	0
0	5	8	0	0	9	8	0
0	4	11	0	1	12	7	0
0	2	14	5	10	12	0	0
0	0	6	13	10	0	0	0




Features




Feature Representation



3rd layer
"Objects"



2nd layer
"Object parts"



1st layer
"Edges"



Pixels

50


Deep Learning

Tensorflow



<https://www.tensorflow.org>


Deep dream generator



<http://deepdreamgenerator.com>




The CIFAR-10 dataset

airplane
automobile
bird
cat
deer
dog
frog
horse
ship
truck



Labeled
training set

With Deep Learning...



horse

car

~ 95% Correct

Is 95% enough?



Tesla Model S. Michael Nagle for The New York Times

Forward-facing camera

Image-processing software can detect lane stripes, signs, stoplights, road signs and other objects.

Forward radar


Reflected microwaves can identify location and speed — but not always the color — of nearby vehicles.

Ultrasonic sensors

Reflected sound waves detect distance to nearby objects.

GPS

Combined with high-precision mapping, GPS determines the car's position on the road.



MAY 2016

<http://wapo.st/2981upr>

“Neither Autopilot nor the driver noticed the white side of the tractor trailer against a brightly lit sky, so the brake was not applied”

Challenges

Viewpoint variation



Scale variation



Deformation



Occlusion



Illumination conditions



Background clutter



Intra-class variation




Text Clustering

Text clustering

IMDB synopses for:

- Robin Hood
- The Matrix
- The King's Speech
- Aladdin
- A Beautiful Mind
- Finding Nemo



CLUSTER 1:

- ~~A?~~ Beautiful Mind
- ~~The~~ Matrix
- ~~The~~ King's Speech


CLUSTER 2:

- ~~Robin~~ Hood
- ~~Ala~~addin
- ~~Finding~~ Nemo

$k = 2$

$k = 2$


Unlabeled data



K-means

$k = 2$


Unlabeled data



K-means

Robin Hood

Told with animals for it's
cast, the story tells of
Robin Hood (a fox) and
Little John (a brown
bear), who rob from the
rich to give to the poor.
[...]




Unlabeled data

$k = 2$

Aladdin
Robin Hood
Finding Nemo

A Beautiful Mind
The Matrix
The King's Speech

K-means




Something simpler...

a) I love CS50. Staff is
awesome, awesome,
awesome!

b) I have a dog and a cat.

c) Best of CS50? Staff.
And cakes. Ok, CS50 staff.

d) My dog keeps chasing my
cat. Dogs!



CLUSTER 1:

a) c)

CLUSTER 2:

b) d)

$k = 2$

$k = 2$

a) I love CS50. Staff is
awesome, awesome,
awesome!


c) Best of CS50? Staff.
And cakes. Ok, CS50 staff.

b) I have a dog and a cat.

d) My dog keeps chasing
my cat. Dogs! |

K-means

a) I love CS50. Staff is
awesome, awesome,
awesome!



a) I love CS50. Staff is
awesome, awesome,
awesome!

a) I love CS50. Staff is awesome, awesome, awesome!

b) I have a dog and a cat.

c) Best of CS50? Staff. And cakes. Ok, CS50 staff.

d) My dog keeps chasing my cat. Dogs! | | |

Bags of words

	awesome	best	cakes	cat	chasing	cs50	dog	dogs	keeps	love	ok	staff
a)	3	0	0	0	0	1	0	0	0	1	0	1
b)	0	0	0	1	0	0	1	0	0	0	0	0
c)	0	1	1	0	0	2	0	0	0	0	1	2
d)	0	0	0	1	1	0	1	1	1	0	0	0

- a) I love CS50. Staff is awesome, awesome, awesome!
- b) I have a dog and a cat.
- c) Best of CS50? Staff. And cakes. Ok, CS50 staff.
- d) My dog keeps chasing my cat. Dogs!

Frequency

	awesome	best	cakes	cat	chasing	cs50	dog	dogs	keeps	love	ok	staff
a)	3/6	0	0	0	0	1/6	0	0	0	1/6	0	1
b)	0	0	0	1/2	0	0	1/2	0	0	0	0	0
c)	0	1/7	1/7	0	0	2/7	0	0	0	0	1/7	2/7
d)	0	0	0	1/5	1/5	0	1/5	1/5	1/5	0	0	0


a) I love CS50. Staff is
awesome, awesome,
awesome!



$(3/6, 0, 0, 0, 0, 1/6, 0, 0, 0, 1/6, 0, 1)$

a) I love CS50. Staff is
awesome, awesome,
awesome!

12 dimensional space



$k = 2$

a) I love CS50. Staff is
awesome, awesome,
awesome!

c) Best of CS50? Staff.
And cakes. Ok, CS50 staff.

b) I have a dog and a cat.



d) My dog keeps chasing
my cat. Dogs! |

K-means

Python code (Unsupervised Learning)

Recap

Handwritten digit classification




0

6

Text clustering

IMDB synopses for:

- Robin Hood
- The Matrix
- The King's Speech
- Aladdin
- A Beautiful Mind
- Finding Nemo



CLUSTER 1:


- A Beautiful Mind
- The Matrix
- The King's Speech

CLUSTER 2:

- Robin Hood
- Aladdin
- Finding Nemo

$k = 2$

Machine Learning?



Search Engines





Image Recognition



Voice Recognition



Natural Language Processing

Machine Learning... so much more

```
  #  #  
 ## ##  
### ###  
#### ####  
#####  
#####  
#####  
#####
```



Machine Learning... so much more



MARCH 2016

“Commentators were convinced [AlphaGo] had made mistakes, but as it racked up wins, **they were forced to concede that perhaps the machine [...] was using strategies its human masters had simply overlooked.**”

*www.economist.com

<https://www.youtube.com/watch?v=JNrXgpSEEIE&t=0m30s>