




# CS 4803 / 7643: Deep Learning

## Topics:

- Image Classification
  - Supervised Learning view
  - K-NN
- 

Dhruv Batra  
Georgia Tech

# Administrativa

- Piazza
  - 165/222 people signed up. Please use that for questions.
- Gradescope/Canvas
  - Anybody not have access?
  - See note on Piazza



# Python+Numpy Tutorial

CS231n Convolutional Neural Networks for Visual Recognition

## Python Numpy Tutorial

This tutorial was contributed by [Justin Johnson](#).

We will use the Python programming language for all assignments in this course. Python is a great general-purpose programming language on its own, but with the help of a few popular libraries (numpy, scipy, matplotlib) it becomes a powerful environment for scientific computing.

We expect that many of you will have some experience with Python and numpy; for the rest of you, this section will serve as a quick crash course both on the Python programming language and on the use of Python for scientific computing.

<http://cs231n.github.io/python-numpy-tutorial/>


# Plan for Today

- Image Classification
  - Supervised Learning view
  - K-NN
- 
- Next time: Linear Classifiers



# Image Classification

# Image Classification: A core task in Computer Vision




[This image](#) by [Nikita](#) is licensed under [CC-BY 2.0](#).

(assume given set of discrete labels)  
{dog, cat, truck, plane, ...}

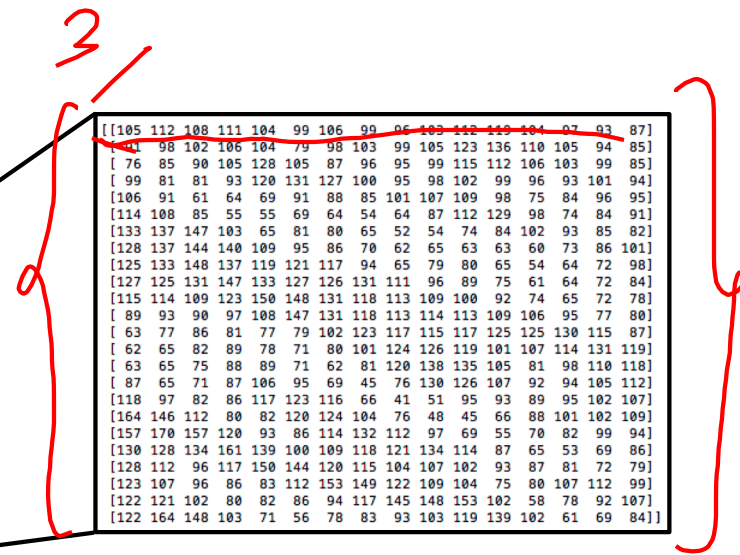


cat

# The Problem: Semantic Gap



This image by Nikita is licensed under CC-BY 2.0




What the computer sees

An image is just a big grid of numbers between [0, 255]:

e.g. 800 x 600 x 3  
(3 channels RGB)


# Challenges: Viewpoint variation



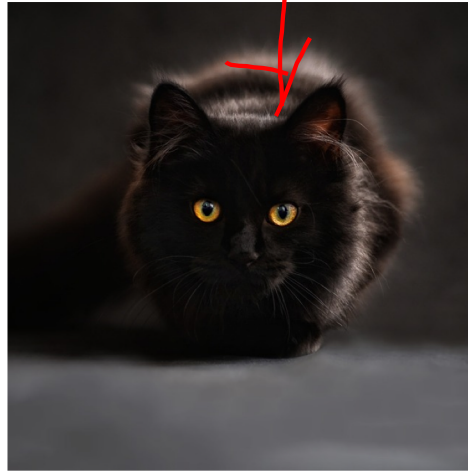
All pixels change when the camera moves!

This image by [Nikita](#) is licensed under [CC-BY 2.0](#)


# Challenges: Illumination



[This image](#) is [CC0 1.0](#) public domain



[This image](#) is [CC0 1.0](#) public domain



[This image](#) is [CC0 1.0](#) public domain



[This image](#) is [CC0 1.0](#) public domain



# Challenges: Deformation




[This image](#) by [Umberto Salvagnin](#) is licensed under [CC-BY 2.0](#)



[This image](#) by [Umberto Salvagnin](#) is licensed under [CC-BY 2.0](#)



[This image](#) by [sare bear](#) is licensed under [CC-BY 2.0](#)



[This image](#) by [Tom Thai](#) is licensed under [CC-BY 2.0](#)



# Challenges: Occlusion



[This image](#) is [CC0 1.0](#) public domain



[This image](#) is [CC0 1.0](#) public domain



[This image](#) by [jonsson](#) is licensed under [CC-BY 2.0](#)

# Challenges: Background Clutter



[This image](#) is [CC0 1.0](#) public domain



[This image](#) is [CC0 1.0](#) public domain

# An image classifier

```
def classify_image(image):  
    # Some magic here?  
    return class_label
```


*list*

*output-list*

Unlike e.g. sorting a list of numbers,

**no obvious way** to hard-code the algorithm for recognizing a cat, or other classes.

# Attempts have been made



John Canny, "A Computational Approach to Edge Detection", IEEE TPAMI 1986



# ML: A Data-Driven Approach

1. Collect a dataset of images and labels
2. Use Machine Learning to train a classifier
3. Evaluate the classifier on new images

Example training set

```
def train(images, labels):  
    # Machine learning!  
    return model
```

```
def predict(model, test_images):  
    # Use model to predict labels  
    return test_labels
```

airplane




automobile



bird



cat



deer



# Notation

Scalars:  $x, y, z, i, \in \mathbb{R}^1$   
 $\vec{x}, \vec{y} \in \mathbb{R}^d$

Matrices  $X, Y$   
R.V, sets

input dim  $d$

output dim  $k$

parameters  $\vec{w}, \vec{\theta} \in \mathbb{R}^d$

#samples  $n, N$

# Supervised Learning

- Input:  $x$   $\in \mathbb{R}^1$   $\mathbb{R}^d$   $\mathbb{R}^{H \times W \times 3}$  (images, text, emails...)  $V = \{1, \dots, d\}$   
 $\{ 'a', 'the', \dots \}$
- Output:  $y$   $\in \{+1, -1\}$  (spam or non-spam...)
- (Unknown) Target Function  $\in \mathbb{R}^k$   
–  $f: X \rightarrow Y$  (the “true” mapping / reality)
- Data  
–  $\{ (x_1, y_1), (x_2, y_2), \dots, (x_N, y_N) \}$   
 $x_i, f(x_i)$

Goal: find  $f$  | Predict  $f(\vec{x})$   
at new  $\vec{x}$

# Supervised Learning

Model Class / Hypothesis Set

$$\underline{H} = \{h: X \rightarrow Y\} \quad \underset{\substack{\uparrow \\ \mathbb{R}^n}}{h(\vec{x})} = y$$

$$\underline{\text{Loss}}(h, D) = \frac{1}{N} \sum_{i=1}^N \underline{L}_i(\underline{h(x_i)}, \underline{y_i})$$

Learning = Search / Opt in MC

$$= \underset{h \in H}{\text{argmin}} \text{Loss}(h, D)$$




# Supervised Learning

- Input:  $x$  (images, text, emails...)
- Output:  $y$  (spam or non-spam...)
- (Unknown) Target Function
  - $f: X \rightarrow Y$  (the “true” mapping / reality)
- Data
  - $\{ (x_1, y_1), (x_2, y_2), \dots, (x_N, y_N) \}$
- Model / Hypothesis Class
  - $H = \{h: X \rightarrow Y\}$
  - e.g.  $y = h(x) = \text{sign}(w^T x)$
- Learning = Search in hypothesis space
  - Find best  $h$  in model class.

# Learning is hard!

## A Learning Problem



#fn

$= 2^d - n$

$2^d - n$

Example	$x_1$	$x_2$	$x_3$	$x_4$	$y$
1	0	0	1	0	0
2	0	1	0	0	0
3	0	0	1	1	1
4	1	0	0	1	1
5	0	1	1	0	0
6	1	1	0	0	0
7	0	1	0	1	0


$x_1$	$x_2$	$x_3$	$x_4$	$y$
0	1	1	1	0, 1

$2^{16}$

# Learning is hard!

- No assumptions = No learning

## A Learning Problem



Example	$x_1$	$x_2$	$x_3$	$x_4$	$y$
1	0	0	1	0	0
2	0	1	0	0	0
3	0	0	1	1	1
4	1	0	0	1	1
5	0	1	1	0	0
6	1	1	0	0	0
7	0	1	0	1	0


# Procedural View

- Training Stage:
  - Training Data  $\{ (x_i, y_i) \} \rightarrow h$  (Learning)
- Testing Stage
  - Test Data  $x \rightarrow h(x)$  (Apply function, Evaluate error)

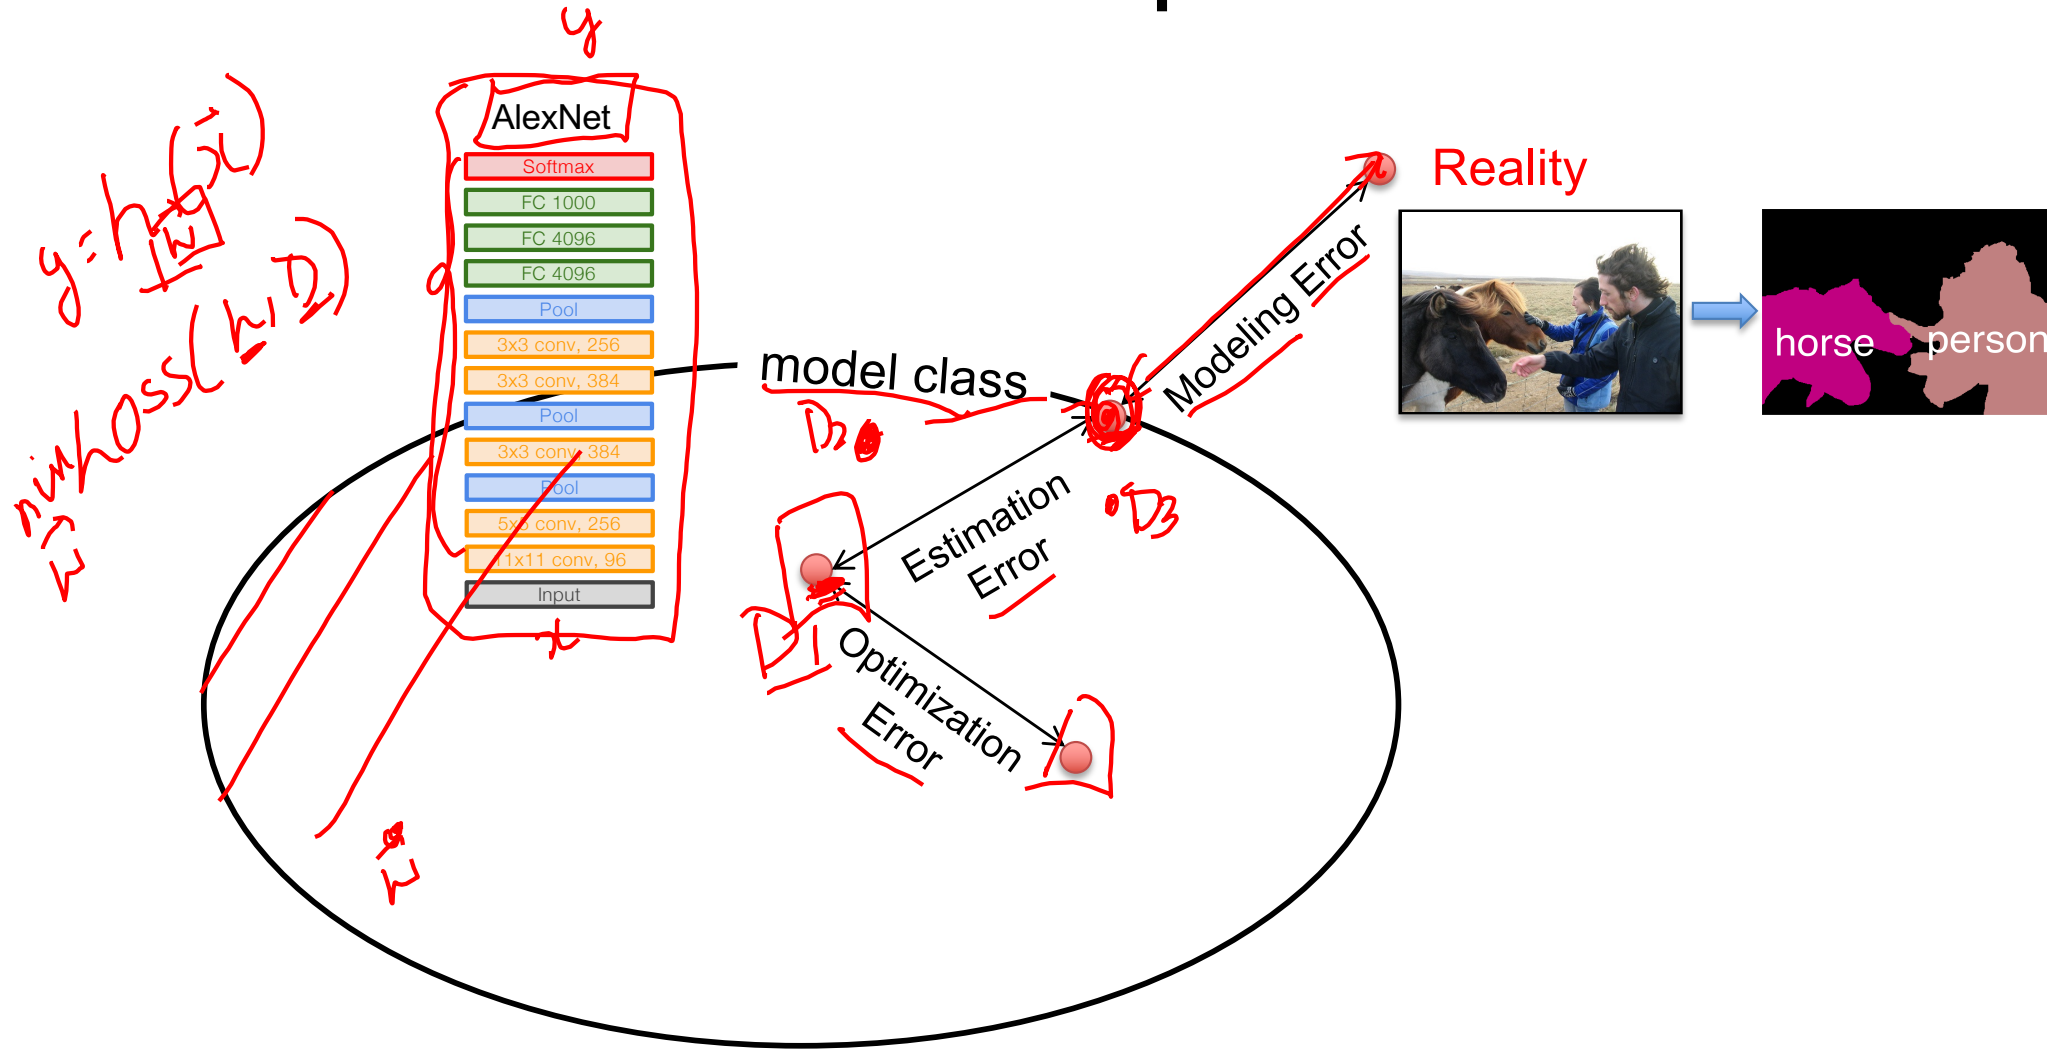
# Statistical Estimation View

- Probabilities to rescue:
  - $X$  and  $Y$  are *random variables*
  - $D = (x_1, y_1), (x_2, y_2), \dots, (x_N, y_N) \sim P(X, Y)$
- IID: Independent Identically Distributed
  - Both training & testing data sampled IID from  $P(X, Y)$
  - Learn on training set
  - Have some hope of *generalizing* to test set

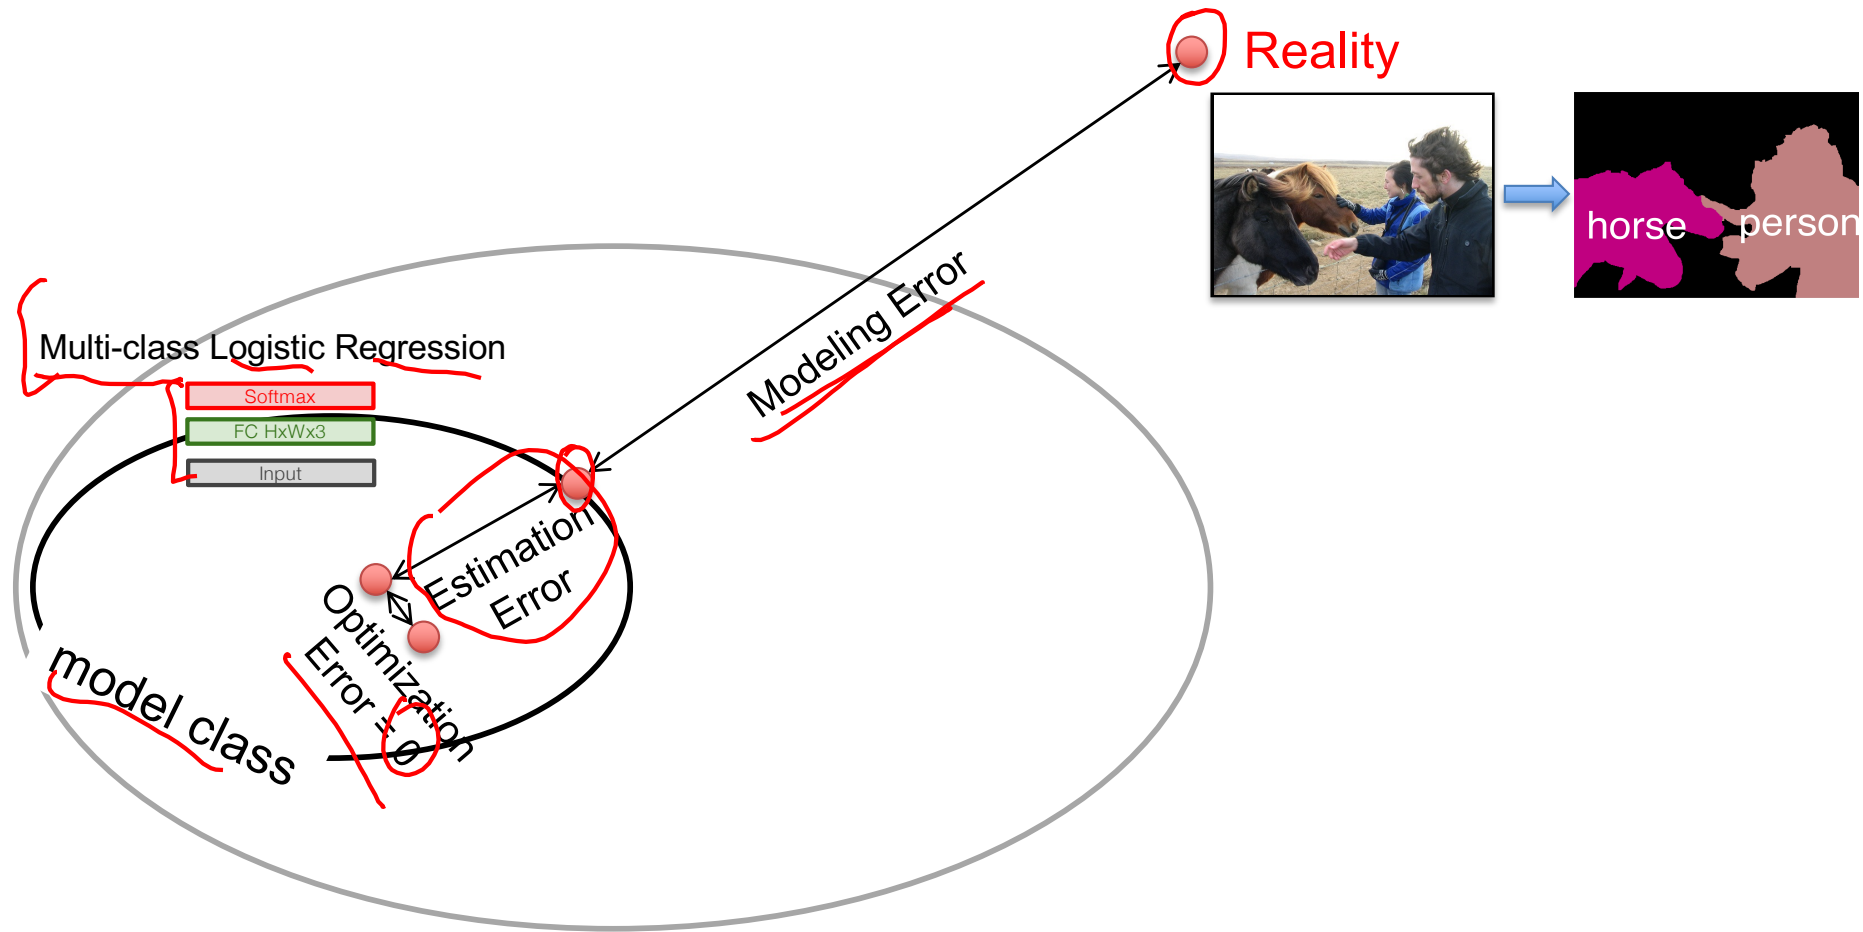
# Error Decomposition



# Error Decomposition




# Error Decomposition







# Error Decomposition

VGG19



model class

Reality



Modeling Error

Estimation Error

Optimization Error

# Error Decomposition

- Approximation/Modeling Error
  - You approximated reality with model
- Estimation Error
  - You tried to learn model with finite data
- Optimization Error
  - You were lazy and couldn't/didn't optimize to completion
- Bayes Error
  - Reality just sucks

# First classifier: Nearest Neighbor

D

```
def train(images, labels):  
    # Machine learning!  
    return model
```



Memorize all  
data and labels

```
def predict(model, test_images):  
    # Use model to predict labels  
    return test_labels
```



Predict the label  
of the most similar  
training image

# Example Dataset: CIFAR10

10 classes

50,000 training images

10,000 testing images

airplane

automobile

bird

cat

deer

dog

frog

horse

ship

truck



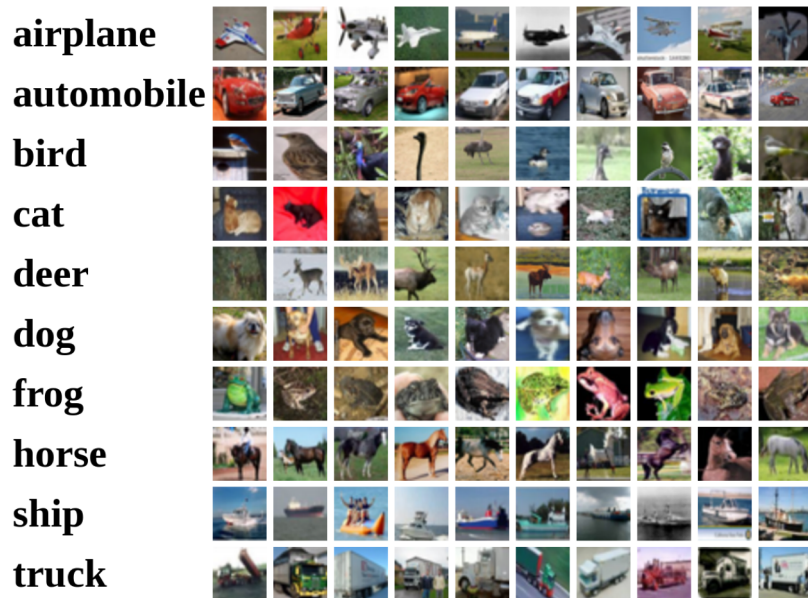
Alex Krizhevsky, "Learning Multiple Layers of Features from Tiny Images", Technical Report, 2009.

# Example Dataset: CIFAR10

10 classes

50,000 training images

10,000 testing images



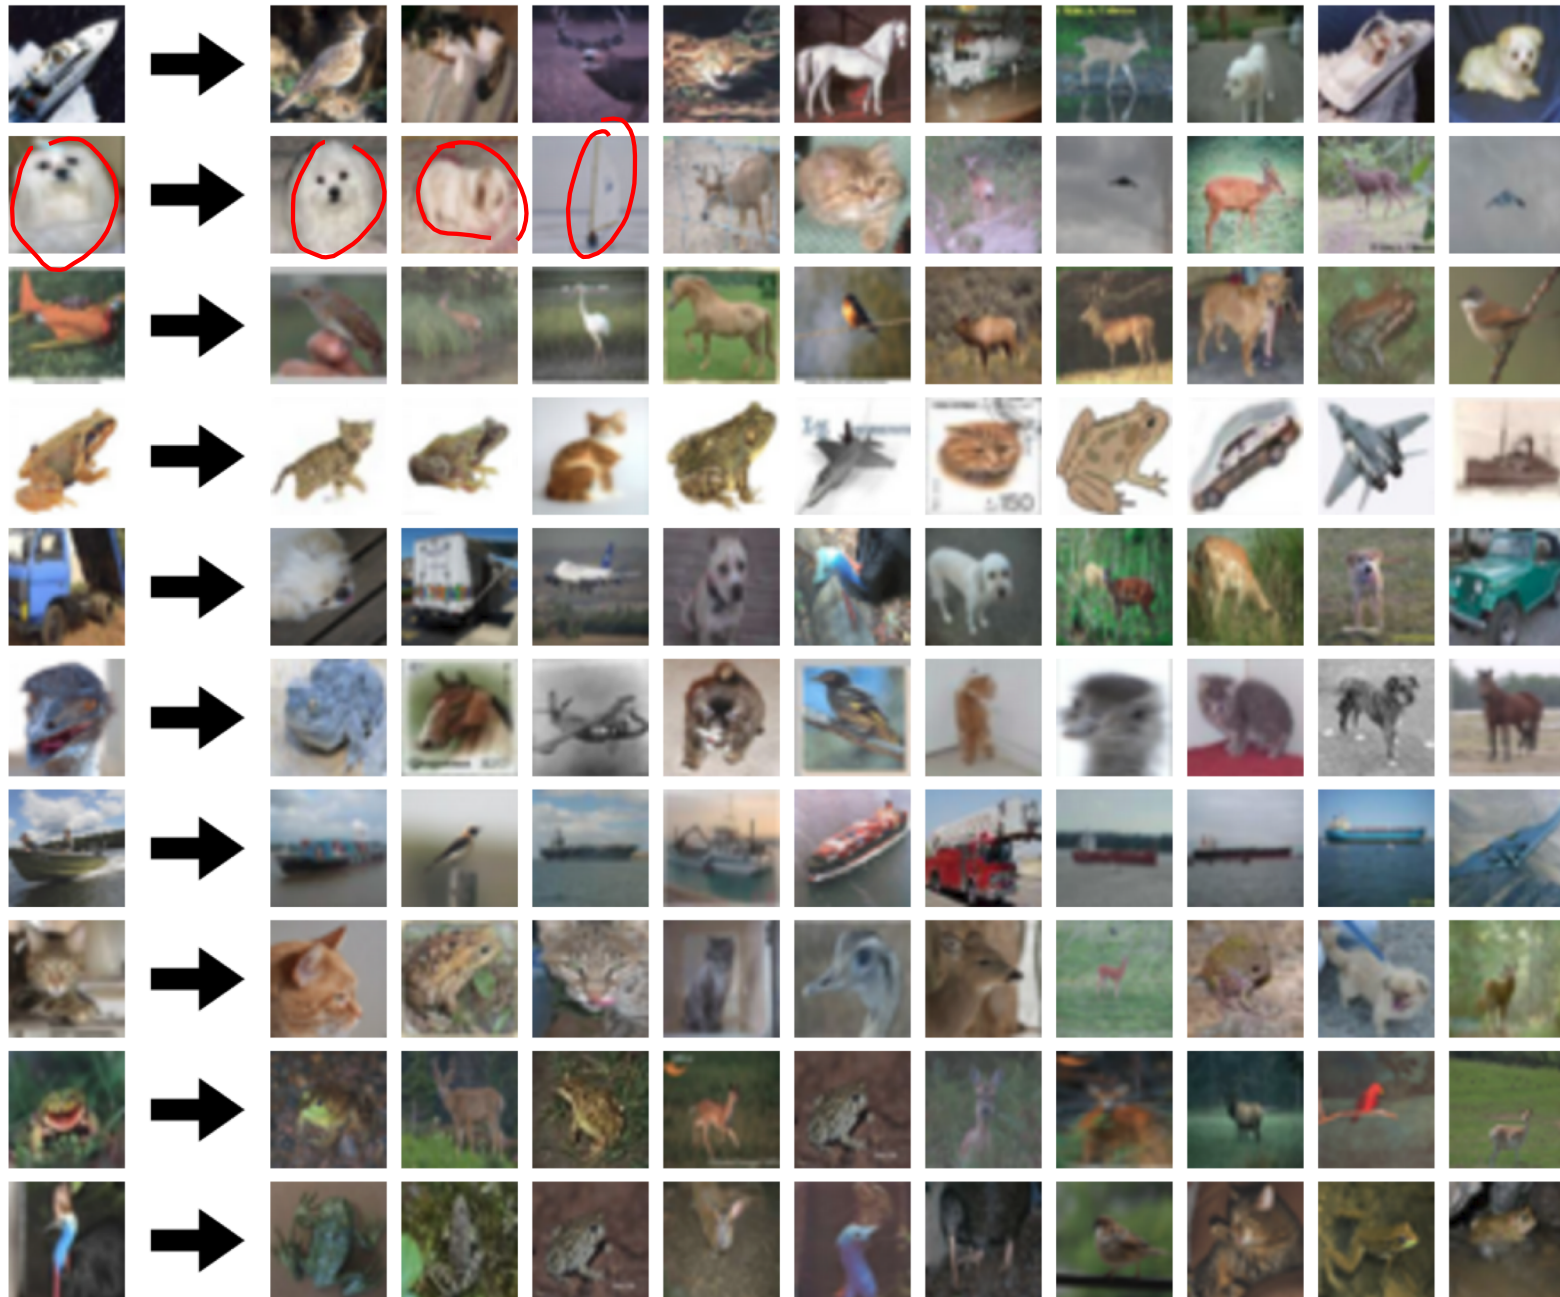
Test images and nearest neighbors



Alex Krizhevsky, "Learning Multiple Layers of Features from Tiny Images", Technical Report, 2009.



# Nearest Neighbours




# Nearest Neighbours



# [Instance/Memory-based Learning]

Four things make a memory based learner:

- A distance metric  $d(\vec{x}_i, \vec{x}_j)$
- How many nearby neighbors to look at?
- A weighting function (optional)
- How to fit with the local points?





# 1-Nearest Neighbour

Four things make a memory based learner:

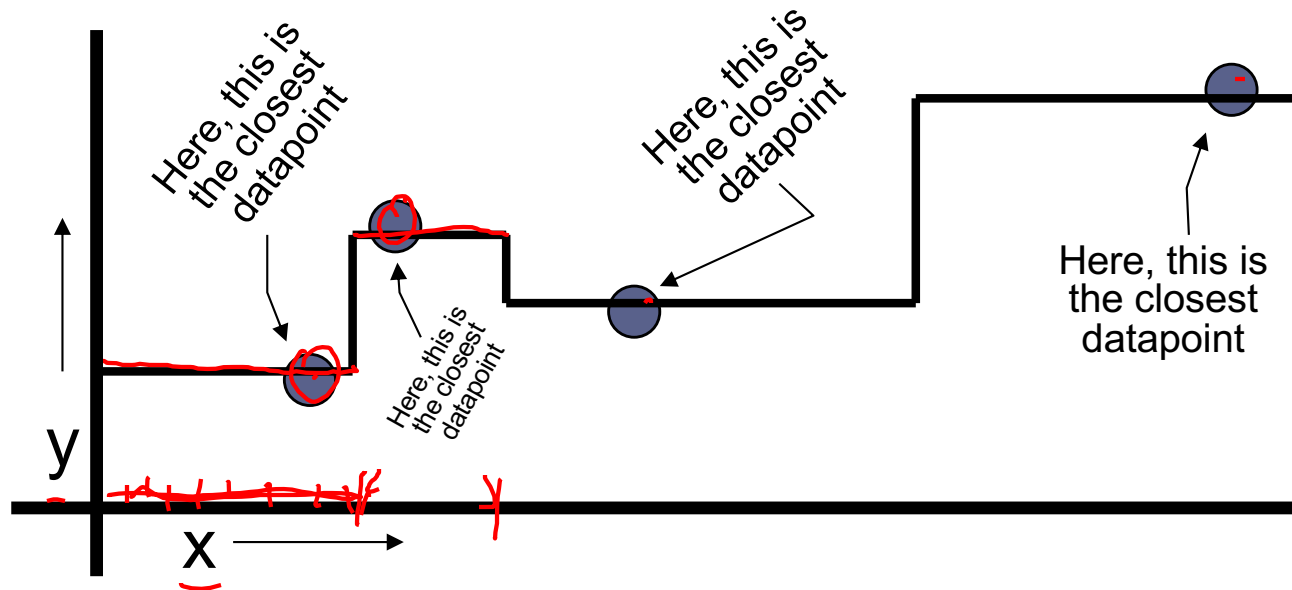
- *A distance metric*
  - Euclidean (and others)
- *How many nearby neighbors to look at?*
  - 1
- *A weighting function (optional)*
  - unused
- *How to fit with the local points?*
  - Just predict the same output as the nearest neighbour.

# k-Nearest Neighbour

Four things make a memory based learner:

- *A distance metric*
  - **Euclidean (and others)**
- *How many nearby neighbors to look at?*
  - **k**
- *A weighting function (optional)*
  - **unused**
- *How to fit with the local points?*
  - **Just predict the average output among the nearest neighbours.**

# 1-NN for Regression



# Distance Metric to compare images

L1 distance:

$$d_1(I_1, I_2) = \sum_P |I_1^P - I_2^P|$$

test image

56	32	10	18
90	23	128	133
24	26	178	200
2	0	255	220

training image

10	20	24	17
8	10	89	100
12	16	178	170
4	32	233	112

pixel-wise absolute value differences

46	12	14	1
82	13	39	33
12	10	0	30
2	32	22	108

add  
→ 456

```

import numpy as np

class NearestNeighbor:
    def __init__(self):
        pass

    def train(self, X, y):
        """ X is N x D where each row is an example. Y is 1-dimension of size N """
        # the nearest neighbor classifier simply remembers all the training data
        self.Xtr = X
        self.ytr = y

    def predict(self, X):
        """ X is N x D where each row is an example we wish to predict label for """
        num_test = X.shape[0]
        # lets make sure that the output type matches the input type
        Ypred = np.zeros(num_test, dtype = self.ytr.dtype)

        # loop over all test rows
        for i in xrange(num_test):
            # find the nearest training image to the i'th test image
            # using the L1 distance (sum of absolute value differences)
            distances = np.sum(np.abs(self.Xtr - X[i,:]), axis = 1)
            min_index = np.argmin(distances) # get the index with smallest distance
            Ypred[i] = self.ytr[min_index] # predict the label of the nearest example

        return Ypred

```

## Nearest Neighbor classifier

```
import numpy as np
```

```
class NearestNeighbor:
```

```
    def __init__(self):  
        pass
```

```
    def train(self, X, y):
```

```
        """ X is N x D where each row is an example. Y is 1-dimension of size N """  
        # the nearest neighbor classifier simply remembers all the training data  
        self.Xtr = X  
        self.ytr = y
```

```
    def predict(self, X):
```

```
        """ X is N x D where each row is an example we wish to predict label for """  
        num_test = X.shape[0]  
        # lets make sure that the output type matches the input type  
        Ypred = np.zeros(num_test, dtype = self.ytr.dtype)  
  
        # loop over all test rows  
        for i in xrange(num_test):  
            # find the nearest training image to the i'th test image  
            # using the L1 distance (sum of absolute value differences)  
            distances = np.sum(np.abs(self.Xtr - X[i,:]), axis = 1)  
            min_index = np.argmin(distances) # get the index with smallest distance  
            Ypred[i] = self.ytr[min_index] # predict the label of the nearest example
```

```
    return Ypred
```

Nearest Neighbor classifier

Memorize training data

```

import numpy as np

class NearestNeighbor:
    def __init__(self):
        pass

    def train(self, X, y):
        """ X is N x D where each row is an example. Y is 1-dimension of size N """
        # the nearest neighbor classifier simply remembers all the training data
        self.Xtr = X
        self.ytr = y

    def predict(self, X):
        """ X is N x D where each row is an example we wish to predict label for """
        num_test = X.shape[0]
        # lets make sure that the output type matches the input type
        Ypred = np.zeros(num_test, dtype = self.ytr.dtype)

        # loop over all test rows
        for i in xrange(num_test):
            # find the nearest training image to the i'th test image
            # using the L1 distance (sum of absolute value differences)
            distances = np.sum(np.abs(self.Xtr - X[i,:]), axis = 1)
            min_index = np.argmin(distances) # get the index with smallest distance
            Ypred[i] = self.ytr[min_index] # predict the label of the nearest example

        return Ypred

```

## Nearest Neighbor classifier

For each test image:  
 Find closest train image  
 Predict label of nearest image



```

import numpy as np

class NearestNeighbor:
    def __init__(self):
        pass

    def train(self, X, y):
        """ X is N x D where each row is an example. Y is 1-dimension of size N """
        # the nearest neighbor classifier simply remembers all the training data
        self.Xtr = X
        self.ytr = y

    def predict(self, X):
        """ X is N x D where each row is an example we wish to predict label for """
        num_test = X.shape[0]
        # lets make sure that the output type matches the input type
        Ypred = np.zeros(num_test, dtype = self.ytr.dtype)

        # loop over all test rows
        for i in xrange(num_test):
            # find the nearest training image to the i'th test image
            # using the L1 distance (sum of absolute value differences)
            distances = np.sum(np.abs(self.Xtr - X[i,:]), axis = 1)
            min_index = np.argmin(distances) # get the index with smallest distance
            Ypred[i] = self.ytr[min_index] # predict the label of the nearest example

        return Ypred

```

## Nearest Neighbor classifier

Q: With N examples, how fast are training and prediction?

```

import numpy as np

class NearestNeighbor:
    def __init__(self):
        pass

    def train(self, X, y):
        """ X is N x D where each row is an example. Y is 1-dimension of size N """
        # the nearest neighbor classifier simply remembers all the training data
        self.Xtr = X
        self.ytr = y

    def predict(self, X):
        """ X is N x D where each row is an example we wish to predict label for """
        num_test = X.shape[0]
        # lets make sure that the output type matches the input type
        Ypred = np.zeros(num_test, dtype = self.ytr.dtype)

        # loop over all test rows
        for i in xrange(num_test):
            # find the nearest training image to the i'th test image
            # using the L1 distance (sum of absolute value differences)
            distances = np.sum(np.abs(self.Xtr - X[i,:]), axis = 1)
            min_index = np.argmin(distances) # get the index with smallest distance
            Ypred[i] = self.ytr[min_index] # predict the label of the nearest example

        return Ypred

```

## Nearest Neighbor classifier

**Q:** With N examples, how fast are training and prediction?

**A:** Train  $O(1)$ ,  
predict  $O(N)$

```

import numpy as np

class NearestNeighbor:
    def __init__(self):
        pass

    def train(self, X, y):
        """ X is N x D where each row is an example. Y is 1-dimension of size N """
        # the nearest neighbor classifier simply remembers all the training data
        self.Xtr = X
        self.ytr = y

    def predict(self, X):
        """ X is N x D where each row is an example we wish to predict label for """
        num_test = X.shape[0]
        # lets make sure that the output type matches the input type
        Ypred = np.zeros(num_test, dtype = self.ytr.dtype)

        # loop over all test rows
        for i in xrange(num_test):
            # find the nearest training image to the i'th test image
            # using the L1 distance (sum of absolute value differences)
            distances = np.sum(np.abs(self.Xtr - X[i,:]), axis = 1)
            min_index = np.argmin(distances) # get the index with smallest distance
            Ypred[i] = self.ytr[min_index] # predict the label of the nearest example

        return Ypred

```


## Nearest Neighbor classifier

**Q:** With N examples, how fast are training and prediction?

**A:** Train  $O(1)$ ,  
predict  $O(N)$

This is bad: we want classifiers that are **fast** at prediction; **slow** for training is ok

# What does this look like?



# Nearest Neighbour

- Demo
  - <http://vision.stanford.edu/teaching/cs231n-demos/knn/>

# Parametric vs Non-Parametric Models

- Does the capacity (size of hypothesis class) grow with size of training data?


- Yes = Non-Parametric Models
- No = Parametric Models

$$H = \{h: X \rightarrow Y\}$$

# K-Nearest Neighbors: Distance Metric


L1 (Manhattan) distance

$$d_1(I_1, I_2) = \sum_p |I_1^p - I_2^p|$$



L2 (Euclidean) distance

$$d_2(I_1, I_2) = \sqrt{\sum_p (I_1^p - I_2^p)^2}$$





# Hyperparameters

What is the best value of k to use?

What is the best distance to use?

These are **hyperparameters**: choices about the algorithm that we set rather than learn

# Hyperparameters

What is the best value of **k** to use?

What is the best **distance** to use?

These are **hyperparameters**: choices about the algorithm that we set rather than learn

Very problem-dependent.

Must try them all out and see what works best.

# Hyperparameters

**Idea #1:** Choose hyperparameters that work best on the data

Your Dataset

# Hyperparameters

**Idea #1:** Choose hyperparameters that work best on the data


**BAD:**  $K = 1$  always works perfectly on training data

Your Dataset


# Hyperparameters

**Idea #1:** Choose hyperparameters that work best on the data

**BAD:**  $K = 1$  always works perfectly on training data



**Idea #2:** Split data into **train** and **test**, choose hyperparameters that work best on test data



# Hyperparameters

**Idea #1:** Choose hyperparameters that work best on the data

**BAD:**  $K = 1$  always works perfectly on training data

Your Dataset

**Idea #2:** Split data into **train** and **test**, choose hyperparameters that work best on test data

**BAD:** No idea how algorithm will perform on new data


train

test

# Hyperparameters

**Idea #1:** Choose hyperparameters that work best on the data

**BAD:**  $K = 1$  always works perfectly on training data



**Idea #2:** Split data into **train** and **test**, choose hyperparameters that work best on test data

**BAD:** No idea how algorithm will perform on new data



**Idea #3:** Split data into **train**, **val**, and **test**; choose hyperparameters on val and evaluate on test

**Better!**



Acc

$K \times$   
 $d^{\otimes}$



# Hyperparameters


Your Dataset

**Idea #4: Cross-Validation:** Split data into **folds**, try each fold as validation and average the results

fold 1	fold 2	fold 3	fold 4	fold 5	test
fold 1	fold 2	fold 3	fold 4	fold 5	test
fold 1	fold 2	fold 3	fold 4	fold 5	test

Useful for small datasets, but not used too frequently in deep learning

# Setting Hyperparameters




Example of  
5-fold cross-validation  
for the value of ~~k~~.

Each point: single  
outcome.

The line goes  
through the mean, bars  
indicated standard  
deviation

(Seems that  $k \approx 7$  works best  
for this data)



# Scene Completion [Hayes & Efros, SIGGRAPH07]












... 200 total

# Context Matching







Graph cut + Poisson blending

Hays and Efros, SIGGRAPH 2007




















# Problems with Instance-Based Learning

- Expensive
  - No Learning: most real work done during testing
  - For every test sample, must search through all dataset – very slow!
  - Must use tricks like approximate nearest neighbour search
- Doesn't work well when large number of irrelevant features
  - Distances overwhelmed by noisy features
- Curse of Dimensionality
  - Distances become meaningless in high dimensions
  - (See proof next)



# k-Nearest Neighbor on images **never used.**

- Very slow at test time
- ~~Distance metrics on pixels are not informative~~




[Original image](#) is  
[CC0 public domain](#)

(all 3 images have same L2 distance to the one on the left)




# k-Nearest Neighbor on images never used.

- Curse of dimensionality



# Curse of Dimensionality


- Consider: Sphere of radius 1 in d-dims



- Consider: an outer  $\epsilon$ -shell in this sphere


- What is  $\frac{\text{shell volume}}{\text{sphere volume}}$ ?

$$\pi r^2$$



$$= \frac{K_d 1^d - K_d (1-\epsilon)^d}{K_d 1^d}$$

$$\frac{4}{3} \pi r^3$$




$$\lim_{d \rightarrow \infty} = \frac{1 - (1-\epsilon)^d}{1} = 1$$

$$K_d r^d$$



# Curse of Dimensionality



# K-Nearest Neighbors: Summary

In **Image classification** we start with a **training set** of images and labels, and must predict labels on the **test set**

The **K-Nearest Neighbors** classifier predicts labels based on nearest training examples

Distance metric and K are **hyperparameters**

Choose hyperparameters using the **validation set**; only run on the test set once at the very end!