

CS 4803 / 7643: Deep Learning

Topics:

- Image Classification
- Supervised Learning view
- K-NN

Dhruv Batra
Georgia Tech

Administrativa

- Waitlist
 - MS-ML or ML-PhD not yet in? Come talk to me.
- Canvas
 - Anybody not have access?
- Piazza
 - 70 people signed up. Please use that for questions.

HW0

- Class Webpage
 - https://www.cc.gatech.edu/classes/AY2019/cs7643_fall/
- Theory
 - https://www.cc.gatech.edu/classes/AY2019/cs7643_fall/assets/hw0.pdf
- Implementation:
 - https://www.cc.gatech.edu/classes/AY2019/cs7643_fall/hw0-a8/

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
- Linear Classifiers



Image Classification

Image Classification: A core task in Computer Vision



Image by [Pete Lin](#) is licensed under [CC BY 4.0](#)

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



cat

The Problem: Semantic Gap




Image by [Justin Johnson](#)
licensed under [CC BY 4.0](#)


```
[[105 112 100 111 104 99 106 99 96 103 112 119 104 97 93 87]
 [ 91 98 102 106 104 79 90 103 99 105 123 136 110 105 94 85]
 [ 76 85 90 105 120 105 87 96 95 99 115 112 106 103 99 85]
 [ 99 81 81 93 120 131 127 100 95 98 102 99 96 93 101 94]
 [106 91 61 64 69 91 88 85 101 107 109 98 75 84 96 95]
 [114 100 85 55 55 69 64 54 64 87 112 129 98 74 84 93]
 [133 137 147 103 65 81 80 65 52 54 74 84 102 93 85 82]
 [120 137 144 140 109 95 86 70 62 65 63 63 60 73 86 101]
 [125 133 140 137 110 121 117 94 65 79 80 65 54 64 72 90]
 [127 125 131 147 133 127 126 131 111 96 89 75 61 64 72 84]
 [115 114 109 123 150 140 131 118 113 109 100 92 74 65 72 70]
 [ 89 93 90 97 100 147 131 118 113 114 113 109 106 95 77 80]
 [ 63 77 86 81 77 79 102 123 117 115 117 125 125 130 115 87]
 [ 62 65 82 89 78 71 80 101 124 126 119 101 107 114 131 119]
 [ 63 65 75 80 80 71 62 81 120 130 135 105 81 90 110 110]
 [ 87 65 71 87 106 95 69 45 76 130 126 107 92 94 105 112]
 [110 97 82 86 117 123 116 66 41 51 95 93 89 95 102 107]
 [104 146 112 80 82 120 124 104 76 40 45 66 80 101 102 109]
 [157 170 157 120 93 86 114 132 112 97 69 55 70 82 99 94]
 [130 120 134 161 139 100 109 118 121 134 114 87 65 53 69 86]
 [120 112 96 117 150 144 120 115 104 107 102 93 87 81 72 79]
 [123 107 96 86 83 112 153 149 122 109 104 75 80 107 112 99]
 [122 121 102 80 82 86 94 117 145 140 153 102 50 70 92 107]
 [122 104 140 103 71 56 70 83 93 103 119 139 102 61 69 84]]
```

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!

Challenges: Illumination



This image is [CC-0](#) public domain



This image is [CC-0](#) public domain



This image is [CC-0](#) public domain



This image is [CC-0](#) public domain

Challenges: Deformation




This image by [Herbert Seligman](#) is licensed under [CC BY 4.0](#)



This image by [Herbert Seligman](#) is licensed under [CC BY 4.0](#)



This image by [mashar](#) is licensed under [CC BY 4.0](#)



This image by [Toni Tini](#) is licensed under [CC BY 4.0](#)

Challenges: Occlusion



This image is [CC0](#) public domain



This image is [CC0](#) public domain



This image by [\[redacted\]](#) is licensed under [CC BY 4.0](#)

Challenges: Background Clutter



This image is [CC0](#) public domain



This image is [CC0](#) public domain


An image classifier

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

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



Supervised Learning

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
 - $h: X \rightarrow Y$
 - $y = h(x) = \text{sign}(w^T x)$
- Learning = Search in hypothesis space
 - Find best h in model class.

Procedural View


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

Statistical Estimation View

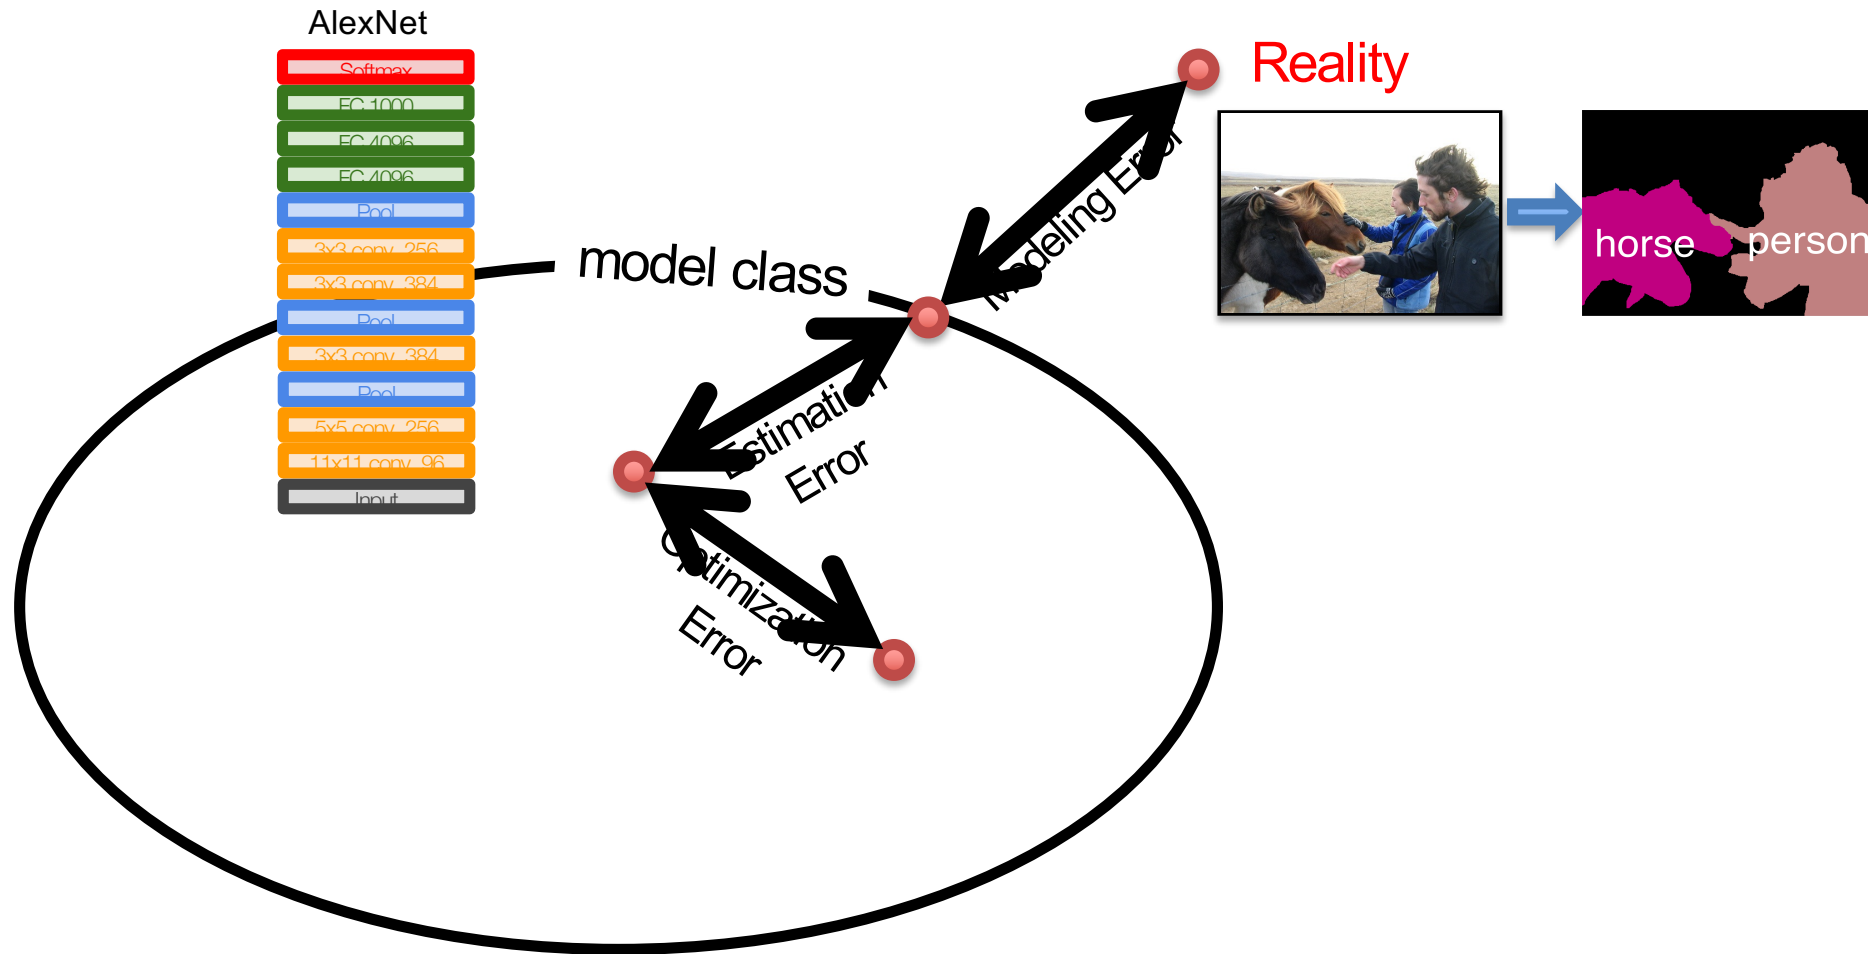
- Probabilities to rescue:
 - X and Y are *random variables*
 - $D = (x_1, y_1), (x_2, y_2), \dots, (x_N, y_N) \sim P(X, Y)$
- IID: Independent Identically Distributed
 - Both training & testing data sampled IID from $P(X, Y)$
 - Learn on training set
 - Have some hope of *generalizing* to test set

Error Decomposition

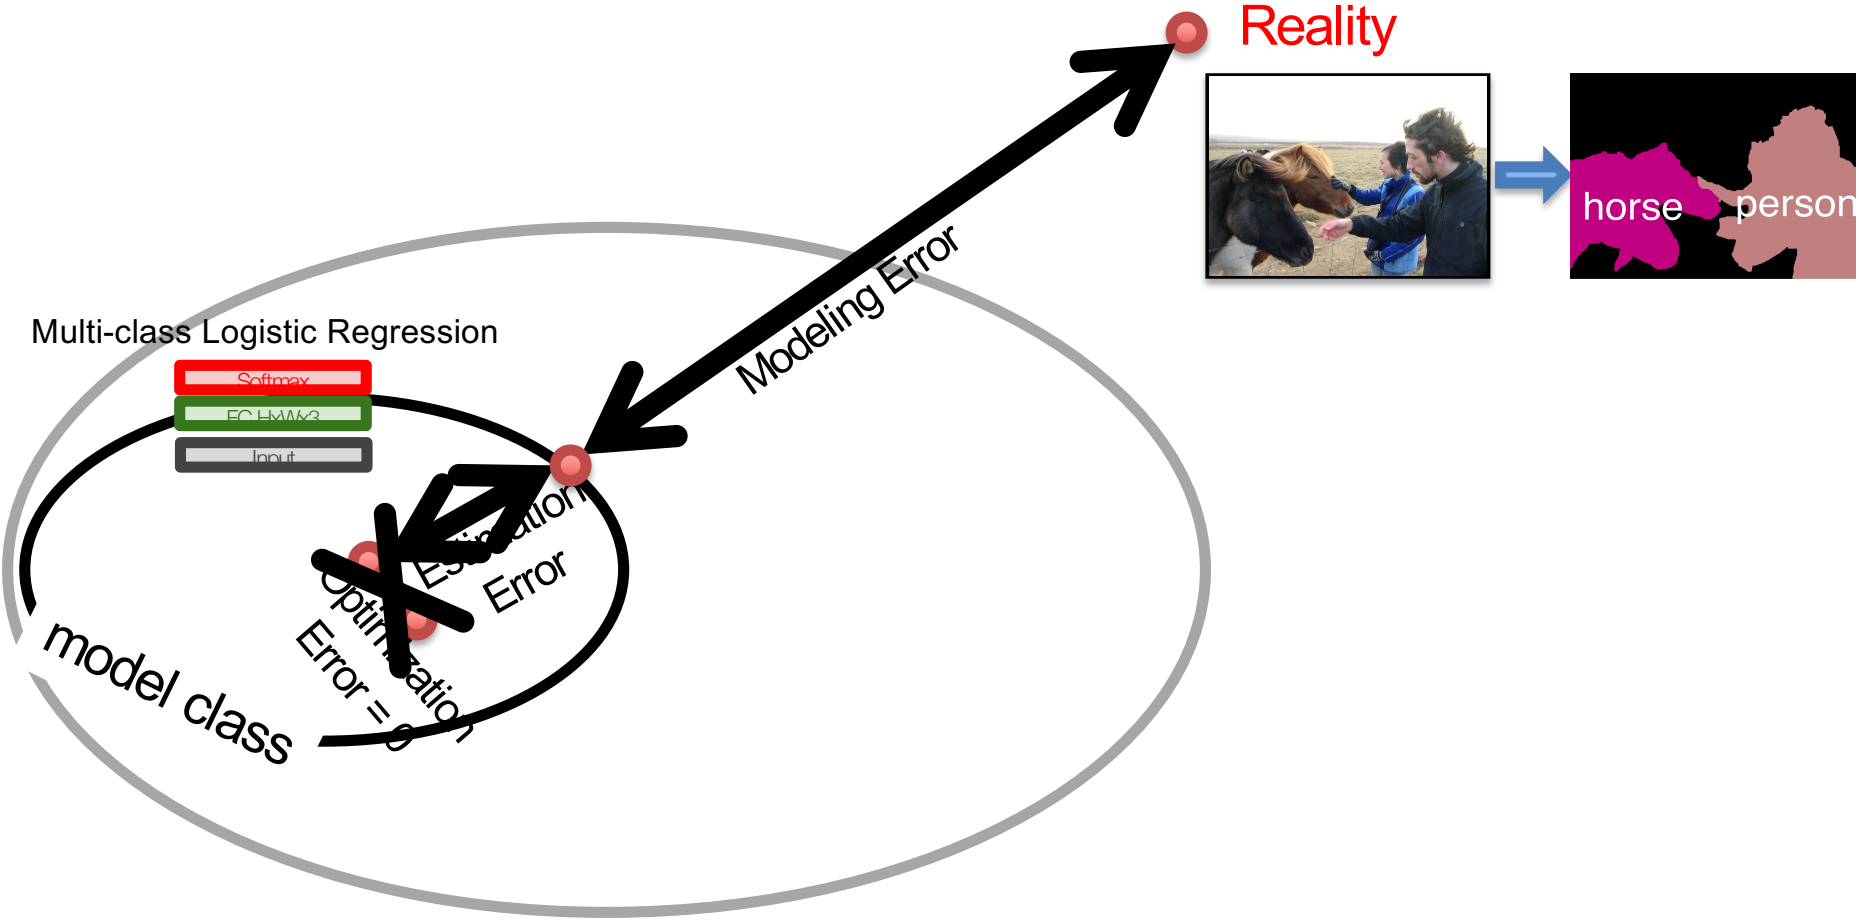
● Reality



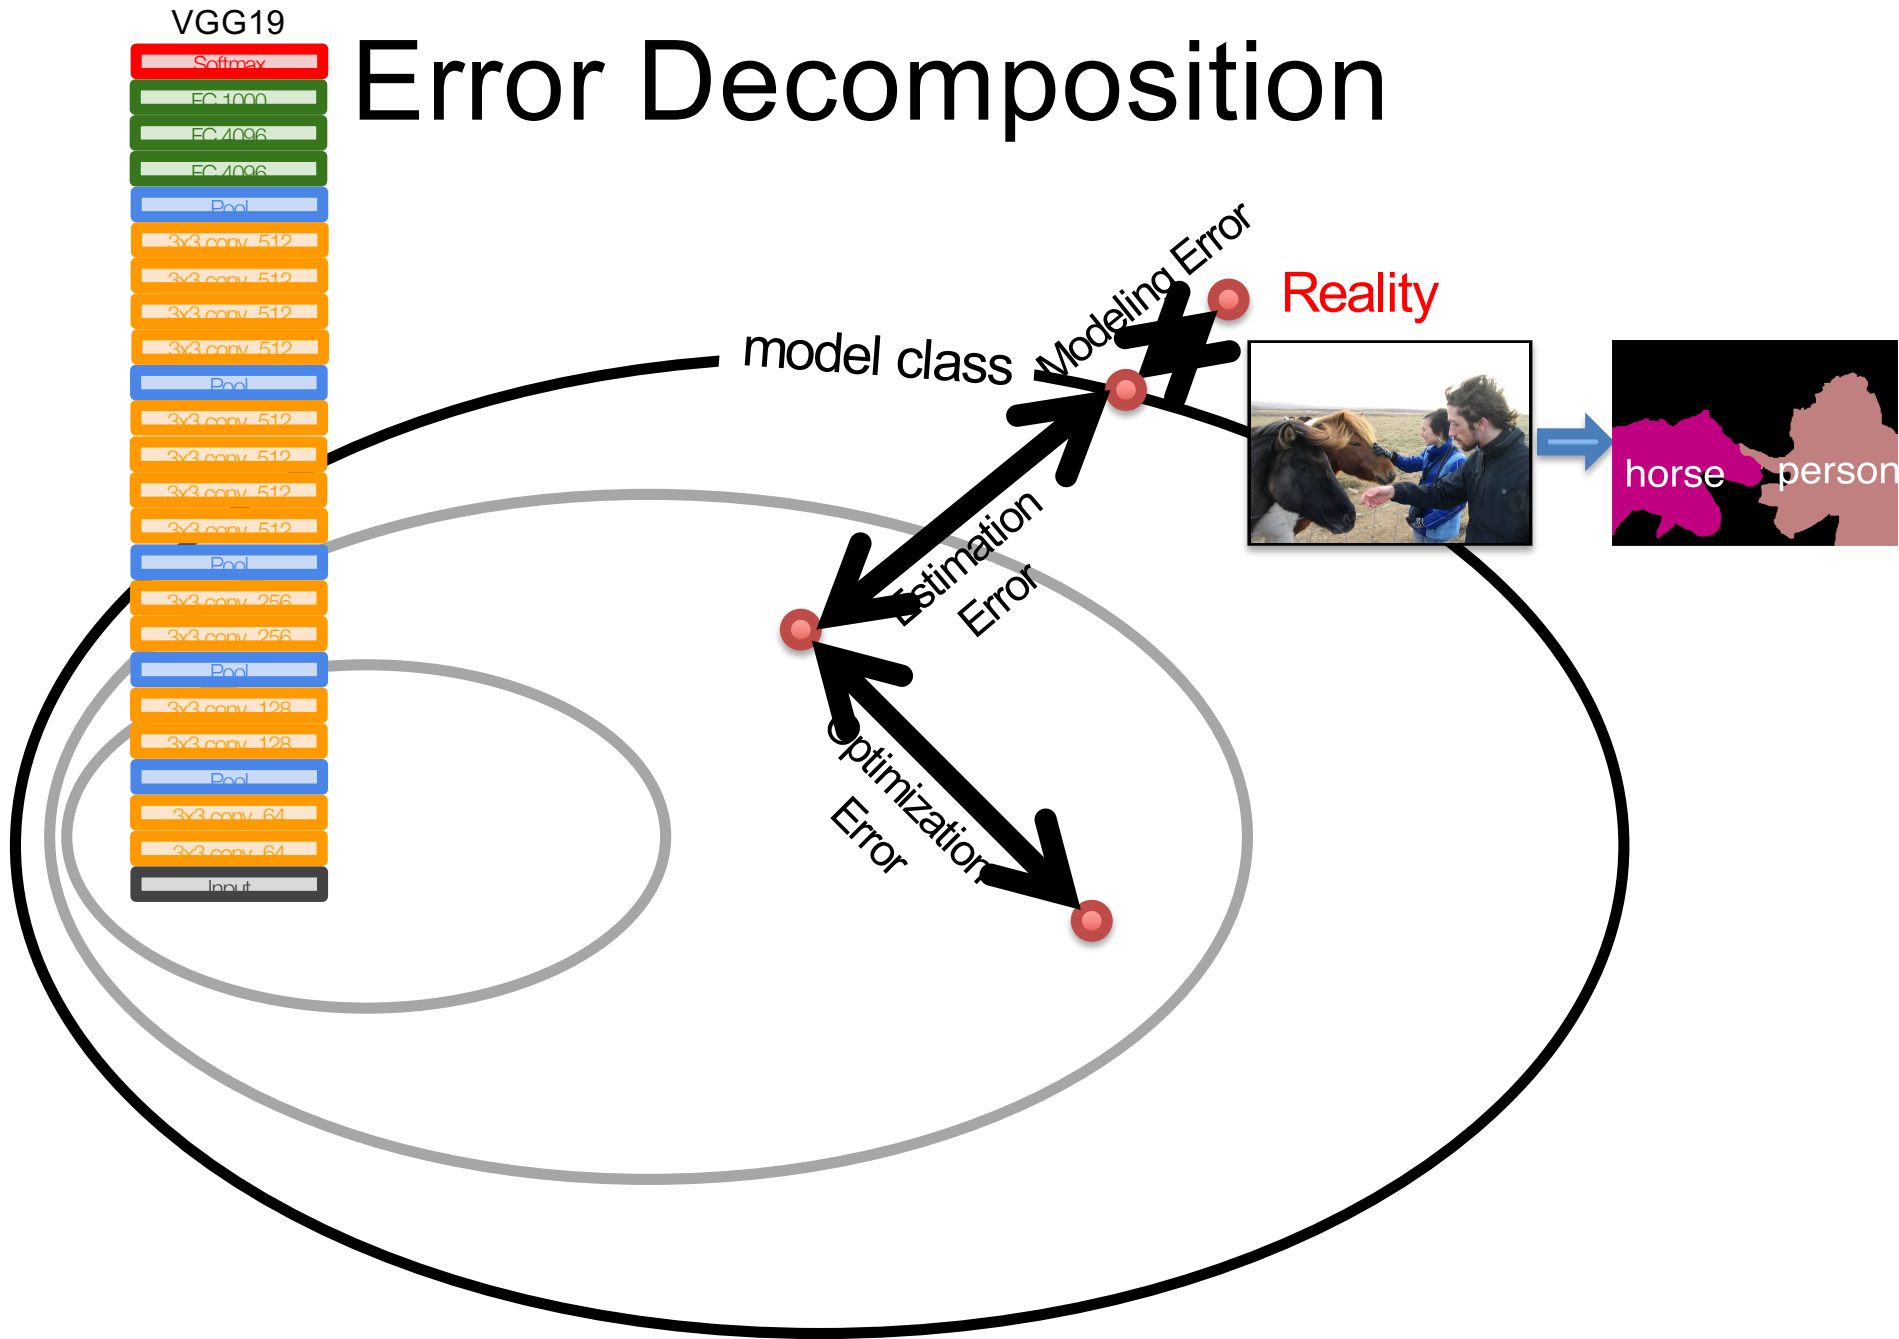
Error Decomposition



Error Decomposition



Error Decomposition



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

Guarantees

- 20 years of research in Learning Theory oversimplified:
- If you have:
 - Enough training data D
 - and H is not too complex
 - then *probably* we can generalize to unseen test data

Learning is hard!

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

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

First classifier: Nearest Neighbor

```
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



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*
- *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				training image				pixel-wise absolute value differences				
56	32	10	18	10	20	24	17	46	12	14	1	add → 456
90	23	128	133	8	10	89	100	82	13	39	33	
24	26	178	200	12	16	178	170	12	10	0	30	
2	0	255	220	4	32	233	112	2	32	22	108	

```

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 1
 - <http://vision.stanford.edu/teaching/cs231n-demos/knn/>
- Demo 2
 - <http://www.cs.technion.ac.il/~rani/LocBoost/>


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

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



Your Dataset

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



train

test

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



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

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

Better!



train


validation

test

Hyperparameters


Your Dataset

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



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 in next lecture)

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

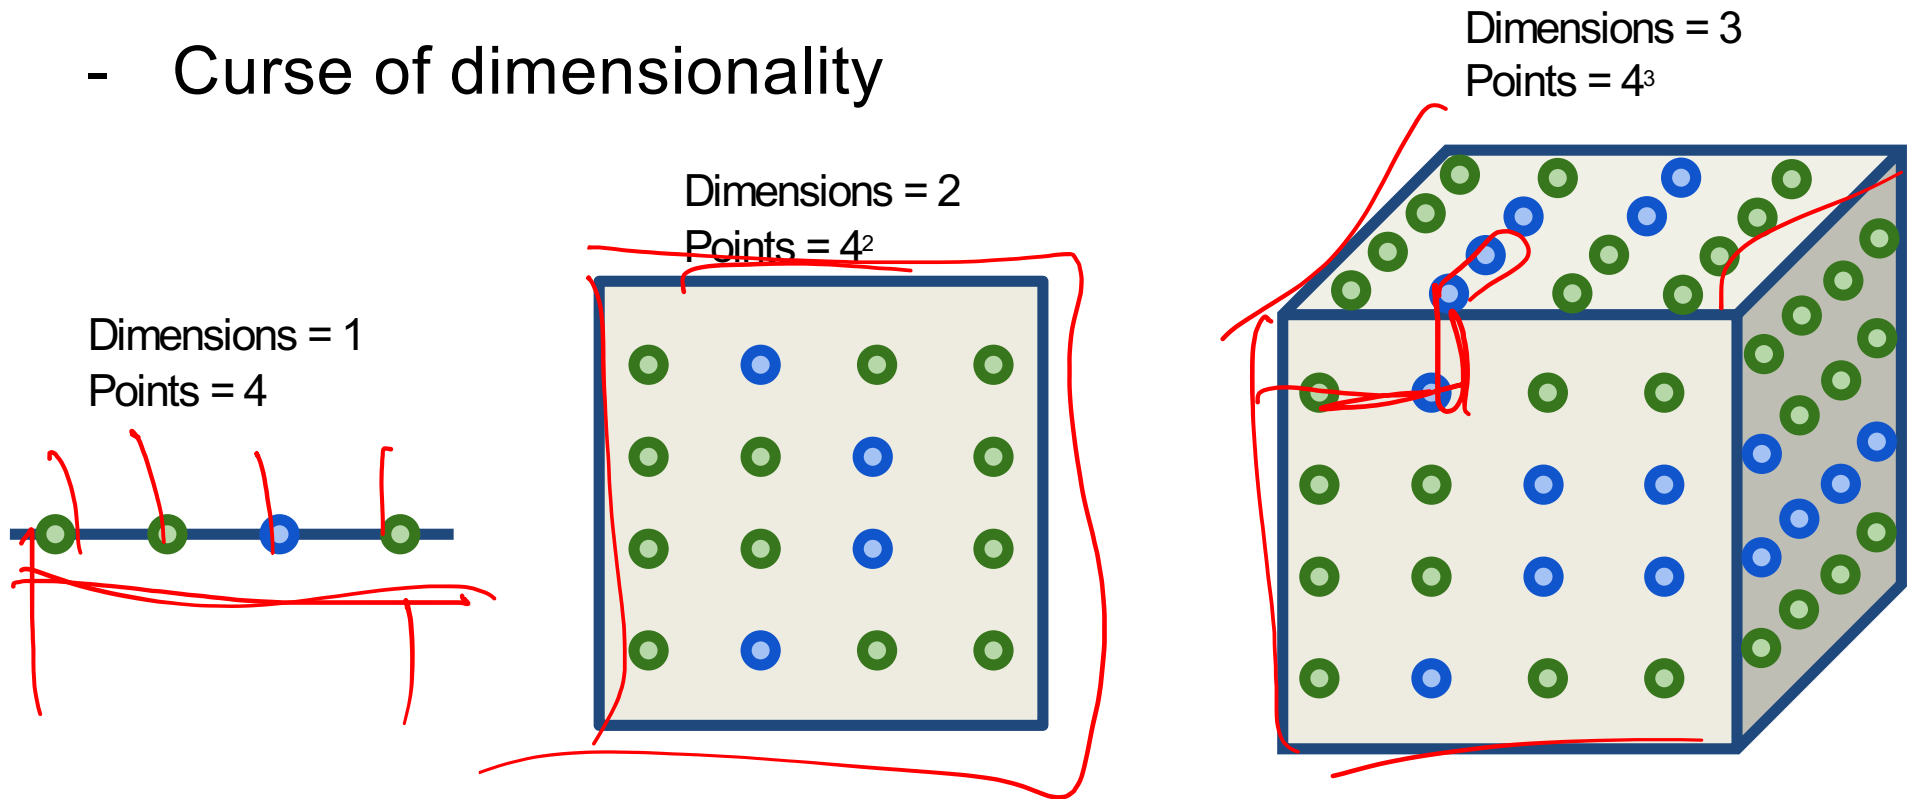
- Very slow at test time
- Distance metrics on pixels are not informative



(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 ϵ -shell in this sphere


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



$$1^d - (1-\epsilon)^d$$

$$1 - (1-\epsilon)^d$$


$$r=1 \quad 2d$$



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

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

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!