

# Lecture 2: Recognition Problems in Vision and k-Nearest Neighbors

CS 444: Deep Learning for Computer Vision

Saurabh Gupta

# Lecture overview


- Different problems in computer vision
- Supervised classification
- Beyond supervised classification: A taxonomy of prediction problems and types of supervision
- Image classification
- k-Nearest Neighbors

# Different Recognition Problems




This image by Nikita is licensed under CC-BY 2.0


**Classification:** Assign image to one of a fixed set of categories



**Object Detection:** Put a bounding box around each instance of a class



**Instance Segmentation:** Mark pixels for each instance of a class




**Semantic Segmentation:** Label each pixels with its category

# Different Recognition Problems




**Image Captioning:** Man riding a horse on a beach



**Depth Prediction:** how far is each pixel in the image




**Keypoint prediction**



**Pose Prediction:** Rotation that aligns object to a canonical pose

# The basic ML framework (for supervised learning)



# The basic ML framework (for supervised learning)

$$y = f(x)$$

↑  
output

↑  
prediction  
function

↙  
input



- **Training (or learning):** given a *training set* of labeled examples  $\{(x_1, y_1), \dots, (x_N, y_N)\}$ , instantiate a predictor  $f$
- **Testing (or inference):** apply  $f$  to a new *test example*  $x$  and output the predicted value  $y = f(x)$
- Rather than hand-defining how 2D projections of apples are different from pears,  $f$  will learn this from the data.


# Is an image classifier all you need?

- Image Classification
- Object Detection
- Instance Segmentation
- Semantic Segmentation
- Image Captioning
- Depth Prediction
- Keypoint Prediction
- Pose Prediction
- ...




# Taxonomy of learning problems

- **Type of output**
  - Classification
  - Regression
    - $y = f(x)$ .  $y$  is an arbitrary scalar and not a class label.
  - Structured prediction
    - $y = f(x)$ .  $y$  is a structured object.



**Depth Prediction:** how far is each pixel in the image



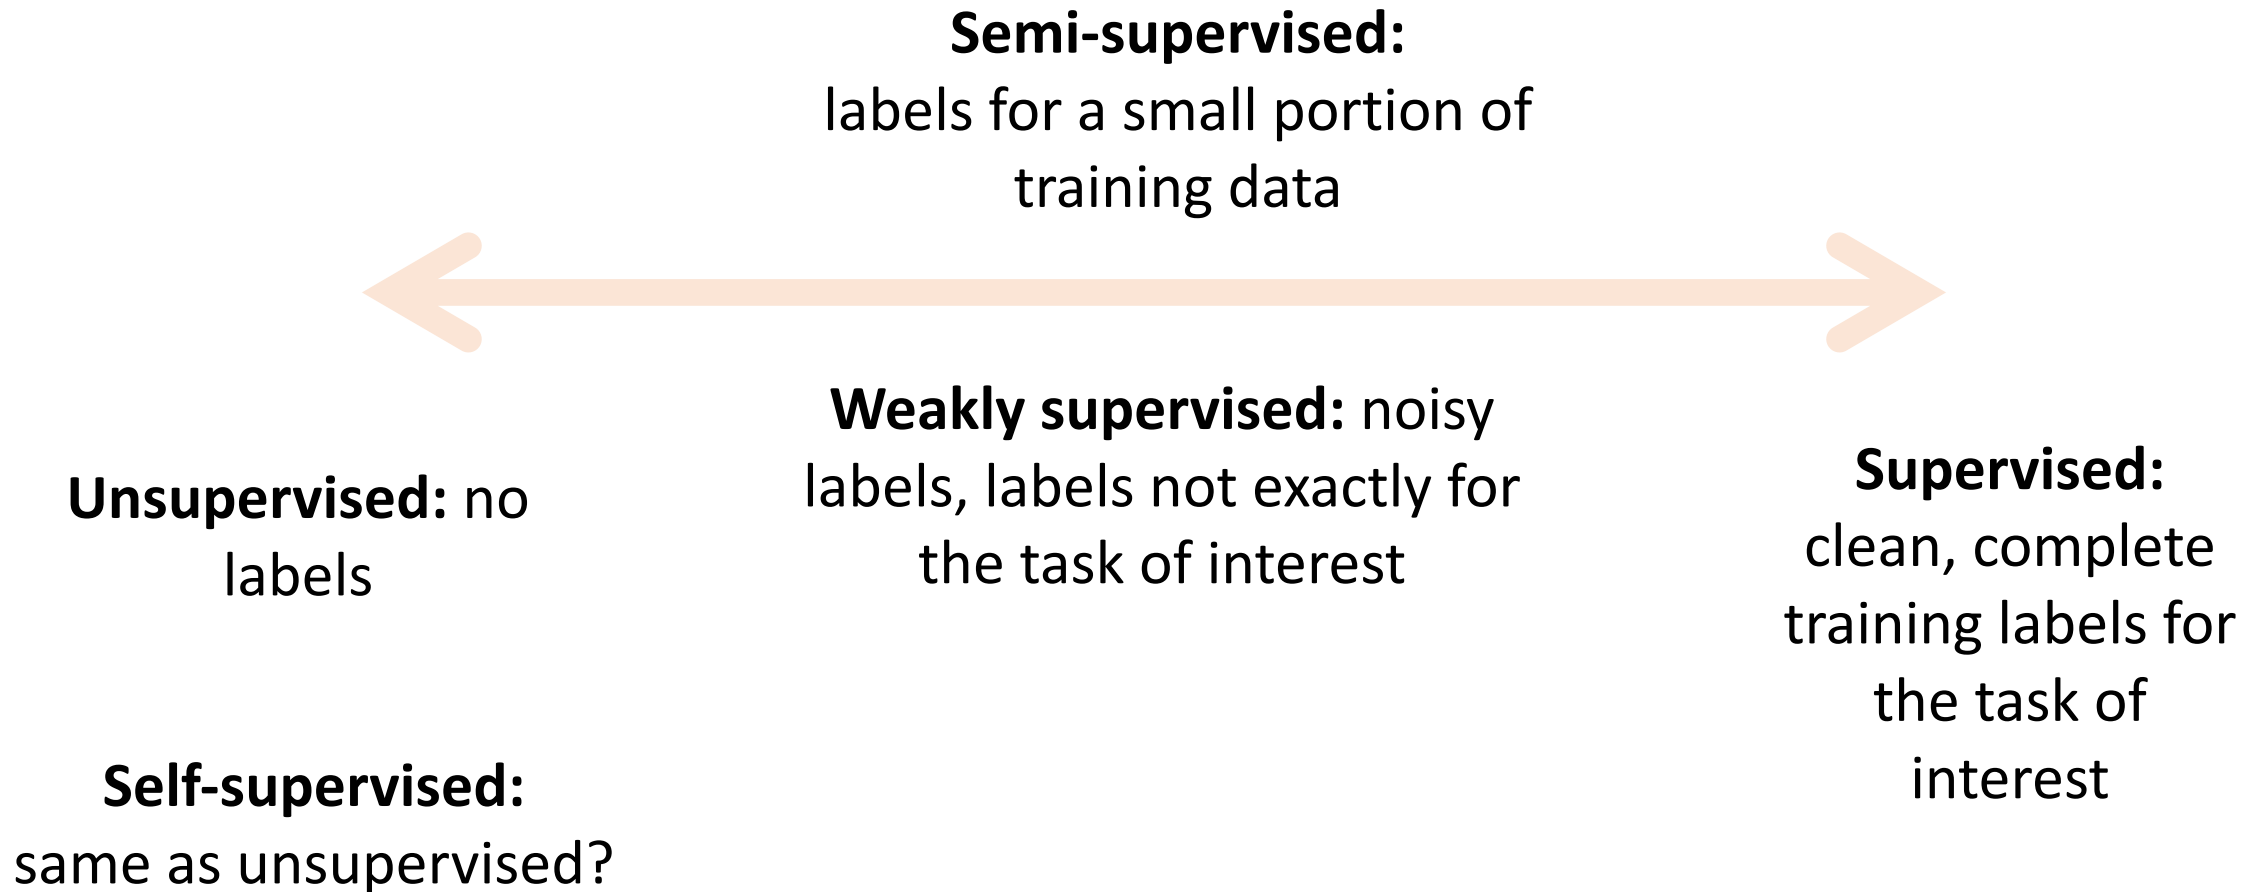
Several computer vision problems have structure in the output space, but often solving a classification problem with some simple post-processing (or even without) ends up being sufficient.



# Taxonomy of learning problems

- **Type of output**
  - Classification
  - Regression
    - $y = f(x)$ .  $y$  is an arbitrary scalar and not a class label.
  - Structured prediction
    - $y = f(x)$ .  $y$  is a structured object.
- **Type of supervision**
  - Supervised learning
  - Unsupervised learning
  - Self-supervised or predictive learning

# Type of supervision

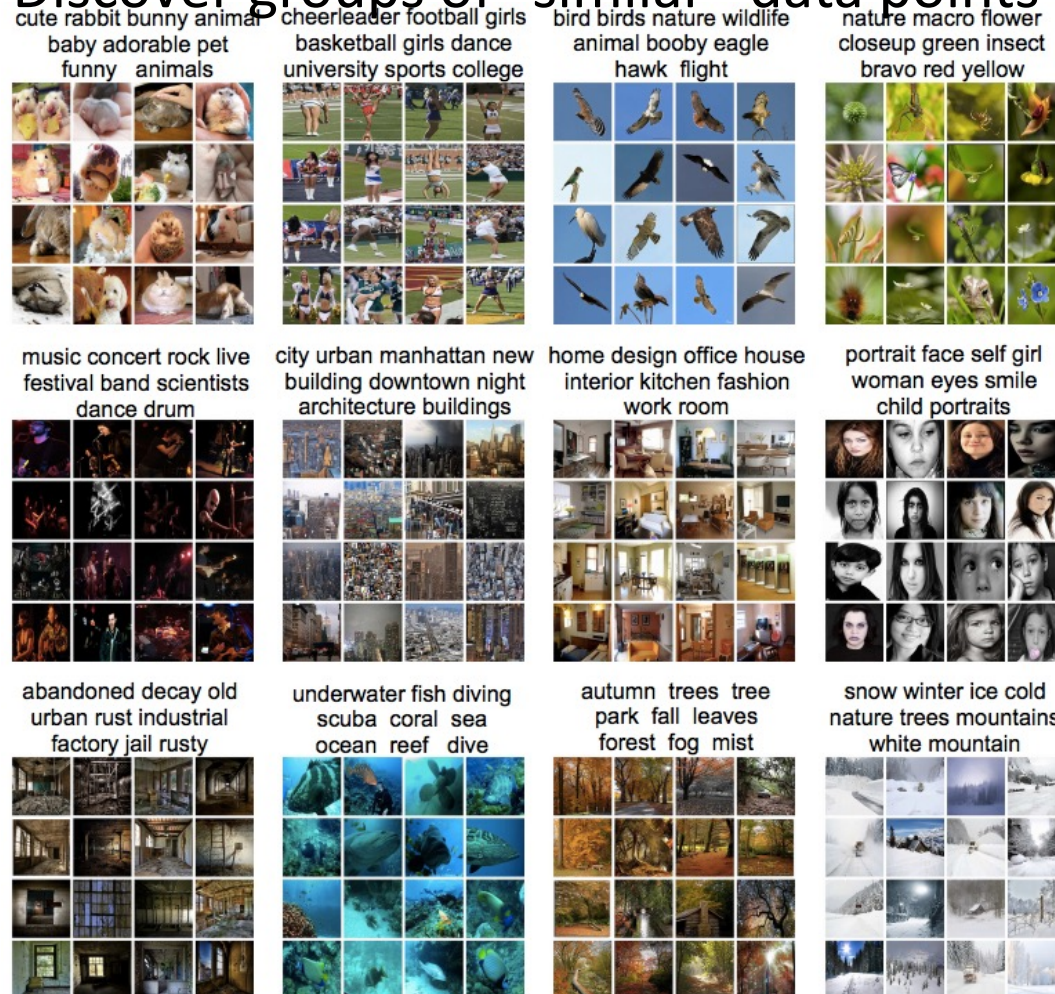


# Unsupervised learning

# Unsupervised learning

- **Clustering**


- Discover groups of “similar” data points



Y. Gong, Q. Ke, M. Isard, and S. Lazebnik. [A Multi-View Embedding Space for Modeling Internet Images, Tags, and Their Semantics](#). IJCV 2014


# Unsupervised learning

- **Dimensionality reduction, manifold learning**
  - Discover a lower-dimensional surface on which the data lives



# Unsupervised learning

- Learning the data distribution
  - **Density estimation:** Find a function that approximates the probability density of the data (i.e., value of the function is high for “typical” points and low for “atypical” points)
  - An extremely hard problem for high-dimensional data...




# Unsupervised learning

- Learning the data distribution
  - **Learning to sample:** Produce samples from a data distribution that mimics the training set

E.g. [Generative adversarial networks](#) (GANs)



“Bedroom”  
(circa 2015)



“Face” (circa  
2015)



4.5 years of GAN progress on face generation.  
[arxiv.org/abs/1406.2661](https://arxiv.org/abs/1406.2661) [arxiv.org/abs/1511.06434](https://arxiv.org/abs/1511.06434)  
[arxiv.org/abs/1606.07536](https://arxiv.org/abs/1606.07536) [arxiv.org/abs/1710.10196](https://arxiv.org/abs/1710.10196)  
[arxiv.org/abs/1812.04948](https://arxiv.org/abs/1812.04948)




6:40 PM · Jan 14, 2019



# Self-supervised or predictive learning

- Use part of the data to predict other parts of the data
  - Example: Masked patch prediction





# Taxonomy of learning problems



- **Type of output**
  - Classification
  - Regression
    - $y = f(x)$ .  $y$  is an arbitrary scalar and not a class label.
  - Structured prediction
    - $y = f(x)$ .  $y$  is a structured object.
- **Type of supervision**
  - Supervised learning
  - Unsupervised learning
  - Self-supervised or predictive learning

# Taxonomy of learning problems

- **Type of output**
  - **Classification**
  - Regression
    - $y = f(x)$ .  $y$  is an arbitrary scalar and not a class label.
  - Structured prediction
    - $y = f(x)$ .  $y$  is a structured object.
- **Type of supervision**
  - **Supervised learning**
  - Unsupervised learning
  - Self-supervised or predictive learning

# Image Classification

# Challenges: Viewpoint Variation



[105	112	108	111	104	99	106	99	96	103	112	119	104	97	93	87]
[ 91	98	102	106	104	79	98	103	99	105	123	136	110	105	94	85]
[ 76	85	90	105	128	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	108	85	55	55	69	64	54	64	87	112	129	98	74	84	91]
[133	137	147	103	65	81	80	65	52	54	74	84	102	93	85	82]
[128	137	144	140	109	95	86	70	62	65	63	60	73	86	101]	
[125	133	148	137	119	121	117	94	65	79	80	65	54	64	72	98]
[127	125	131	147	133	127	126	131	111	96	89	75	61	64	72	84]
[115	114	109	123	150	148	131	118	113	109	100	92	74	65	72	78]
[ 89	93	90	97	108	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	88	89	71	62	81	120	138	135	105	81	98	110	118]
[ 87	65	71	87	106	95	69	45	76	130	126	107	92	94	105	112]
[118	97	82	86	117	123	116	66	41	51	95	93	89	95	102	107]
[164	146	112	80	82	120	124	104	76	48	45	66	88	101	102	109]
[157	170	157	120	93	86	114	132	112	97	69	55	70	82	99	94]
[130	128	134	161	139	100	109	118	121	134	114	87	65	53	69	86]
[128	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	148	153	102	58	78	92	107]
[122	164	148	103	71	56	78	83	93	103	119	139	102	61	69	84]

All pixels change when the camera moves!

# Challenges: Intraclass Variation



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

# Challenges: Fine-Grained Categories

Maine Coon



[This image is free for use under the Pixabay License](#)

Ragdoll



[This image is CC0 public domain](#)

American Shorthair



[This image is CC0 public domain](#)

# Challenges: Background Clutter



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



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

# Challenges: Illumination Changes




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 [jonsso](#) is licensed under [CC-BY 2.0](#)

# Image Classification Datasets: MNIST



**10 classes:** Digits 0 to 9  
**28x28** grayscale images  
**50k** training images  
**10k** test images

# Image Classification Datasets: MNIST



**10 classes:** Digits 0 to 9  
**28x28** grayscale images  
**50k** training images  
**10k** test images

“Drosophila of computer vision”

Results from MNIST often do not hold on more complex datasets!

# Image Classification Datasets: CIFAR10


**airplane**




**automobile**




**bird**



**cat**



**deer**




**dog**



**frog**




**horse**



**ship**



**truck**




**10 classes**

**50k training images (5k per class)**

**10k testing images (1k per class)**

**32x32 RGB images**

# Image Classification Datasets: CIFAR100



**100** classes

**50k** training images (500 per class)

**10k** testing images (100 per class)

**32x32 RGB** images

**20 superclasses** with 5 classes each:

Aquatic mammals: beaver, dolphin, otter, seal, whale

Trees: Maple, oak, palm, pine, willow

# Image Classification Datasets: ImageNet

**1000** classes

**~1.3M** training images (~1.3K per class)

**50K** validation images (50 per class)

**100K** test images (100 per class)

Performance metric: **Top 5 accuracy**


Algorithm predicts 5 labels for each image; one of them needs to be right




flamingo



cock



ruffed grouse




quail




partridge


...




Egyptian cat



Persian cat



Siamese cat




tabby



lynx


...




dalmatian




keeshond



miniature schnauzer



standard schnauzer



giant schnauzer

# Image Classification Datasets: ImageNet


**1000** classes

**~1.3M** training images (~1.3K per class)

**50K** validation images (50 per class)

**100K** test images (100 per class)

test labels are secret!



flamingo

cock

ruffed grouse

quail

partridge

...



Egyptian cat

Persian cat

Siamese cat

tabby

lynx

...



dalmatian

keeshond

miniature schnauzer

standard schnauzer

giant schnauzer

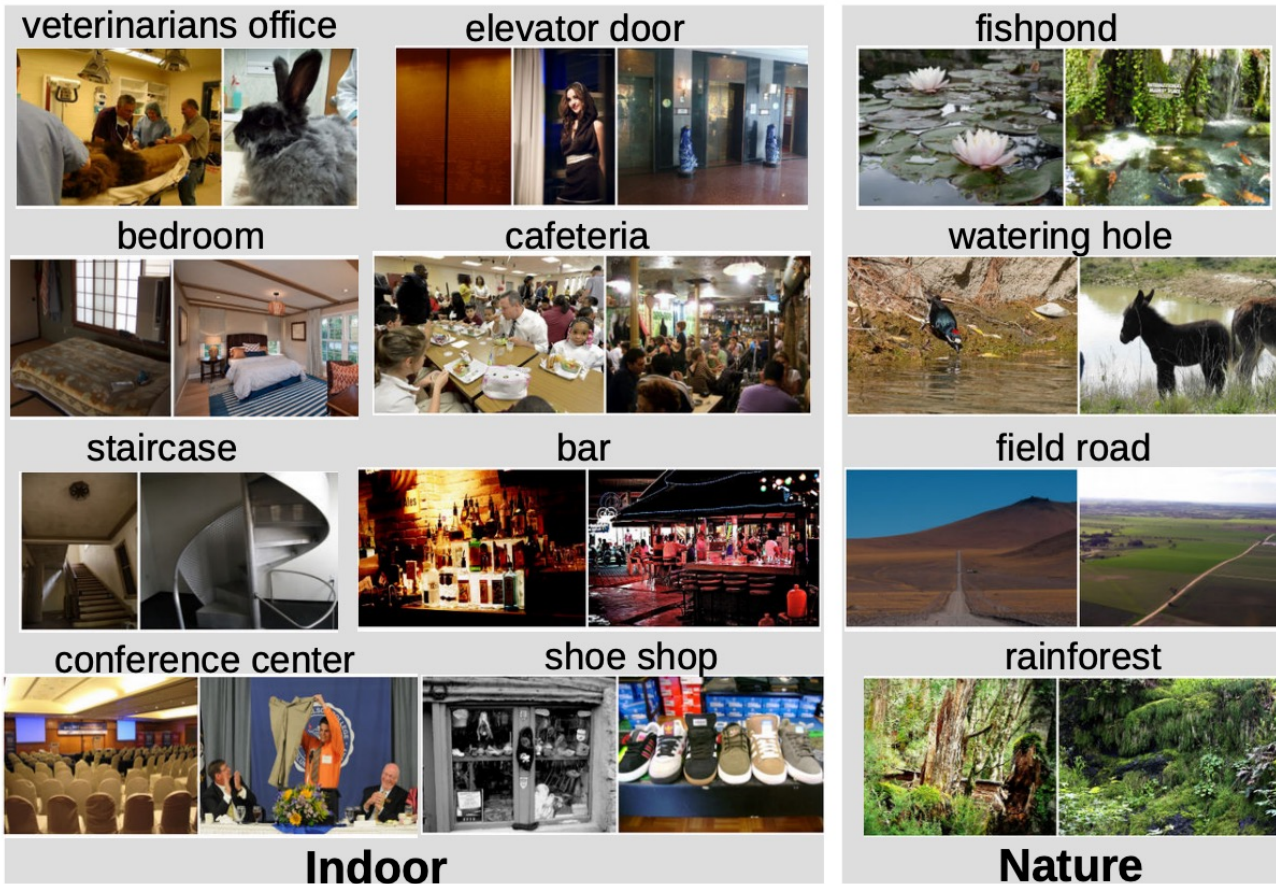
Images have variable size, but often  
resized to **256x256** for training

There is also a 22k category version of  
ImageNet, but less commonly used

Deng et al, "ImageNet: A Large-Scale Hierarchical Image Database", CVPR 2009  
Russakovsky et al, "ImageNet Large Scale Visual Recognition Challenge", IJCV 2015



# Image Classification Datasets: MIT Places



**365 classes** of different scene types

~**8M** training images

**18.25K** val images (50 per class)

**328.5K** test images (900 per class)

Images have variable size, often  
resize to **256x256** for training

Zhou et al, "Places: A 10 million Image Database for Scene Recognition", TPAMI 2017

# k-Nearest Neighbors

# 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

# 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

# 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

```
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
```

For each test image:  
Find nearest training image  
Return 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 is training?



```

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 is training?

**A:**  $O(1)$

```

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 is training?

**A:**  $O(1)$

**Q:** With N examples,  
how fast is testing?

```

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 is training?

**A:**  $O(1)$

**Q:** With N examples,  
how fast is testing?

**A:**  $O(N)$

```

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 is training?

**A:**  $O(1)$

**Q:** With N examples,  
how fast is testing?

**A:**  $O(N)$

This is **bad**: We can  
afford slow training, but  
we need fast testing!

# 1-Nearest Neighbor Complexity

Method	Query Time Complexity
<i>Exact Algorithms:</i>	
RP Tree	$O((d' \log d')^{d'} + \log n)$
Spill Tree	$O(d'^{d'} + \log n)$
Karger & Ruhl (2002)	$O(2^{3d'} \log n)$
Navigating Net	$2^{O(d')} \log n$
Cover Tree	$O(2^{12d'} \log n)$
Rank Cover Tree	$O(2^{O(d' \log h)} n^{2/h})$ for $h \geq 3$
DCI	$O(d \max(\log n, n^{1-1/d'}))$
Prioritized DCI	$O(d \max(\log n, n^{1-m/d'}))$
(Proposed Method)	$+(m \log m) \max(\log n, n^{1-1/d'})$ for $m \geq 1$
<i>Approximate Algorithms:</i>	
$k$ -d Tree	$O((1/\epsilon)^d \log n)$
BBD Tree	$O((6/\epsilon)^d \log n)$
LSH	$\approx O(dn^{1/(1+\epsilon)^2})$

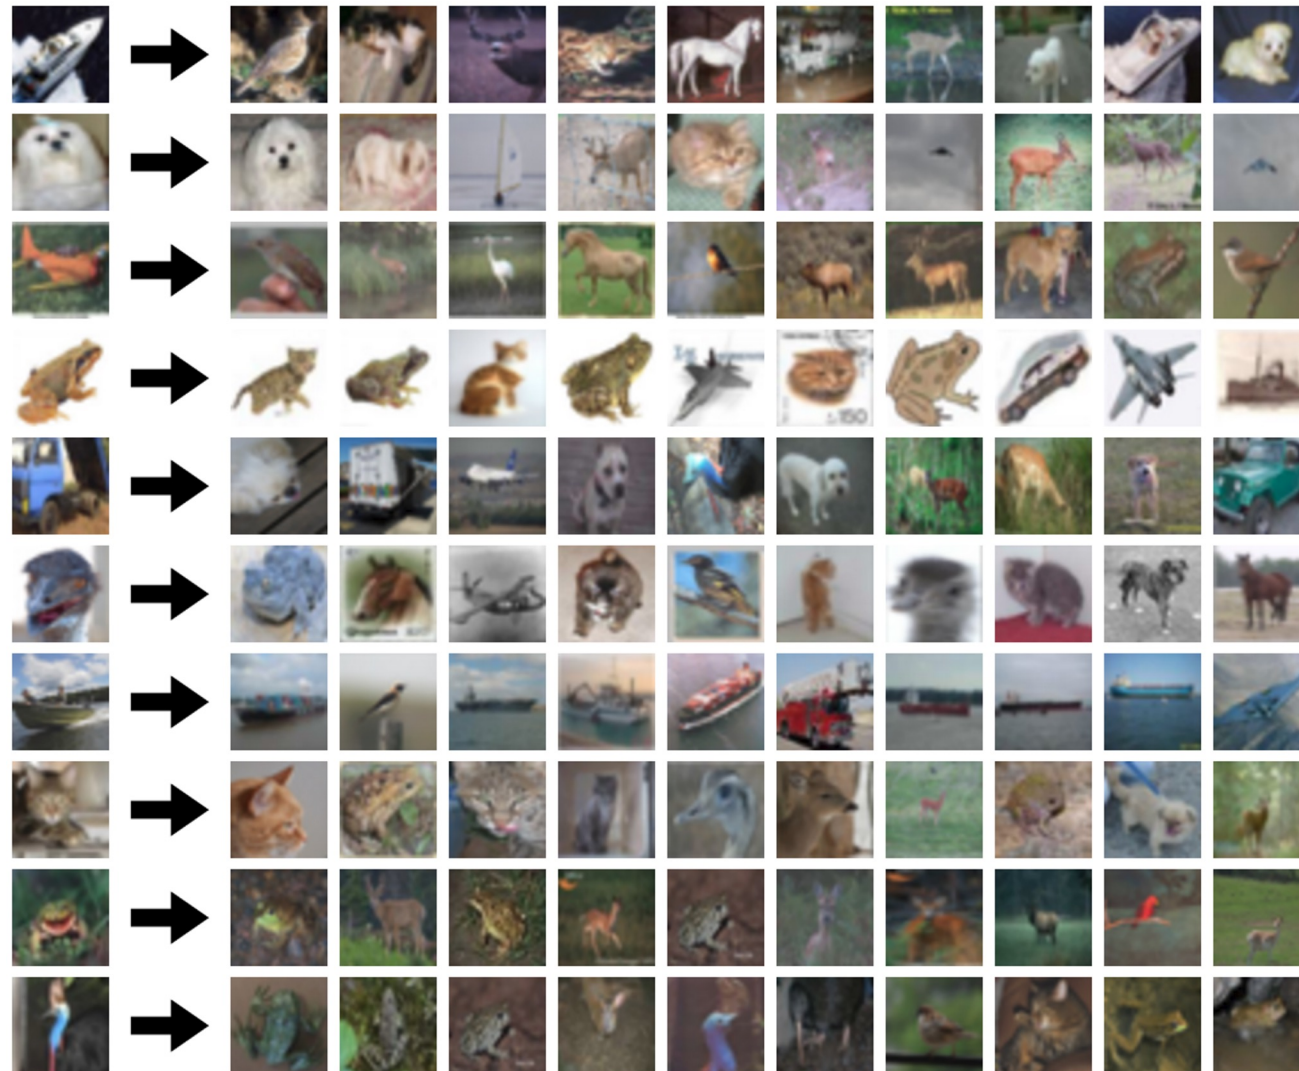
Table 1. Query time complexities of various algorithms for 1-NN search. Ambient dimensionality, intrinsic dimensionality, dataset size and approximation ratio are denoted as  $d$ ,  $d'$ ,  $n$  and  $1 + \epsilon$ . A visualization of the growth of various time complexities as a function of the intrinsic dimensionality is shown in Figure 1.

**Bad news overall:**  
Exponential in dimensionality or  
(almost) linear in number of data points.

Good Implementation:

<https://github.com/facebookresearch/faiss>


# What does this look like?



# What does this look like?




# Nearest Neighbor Decision Boundaries






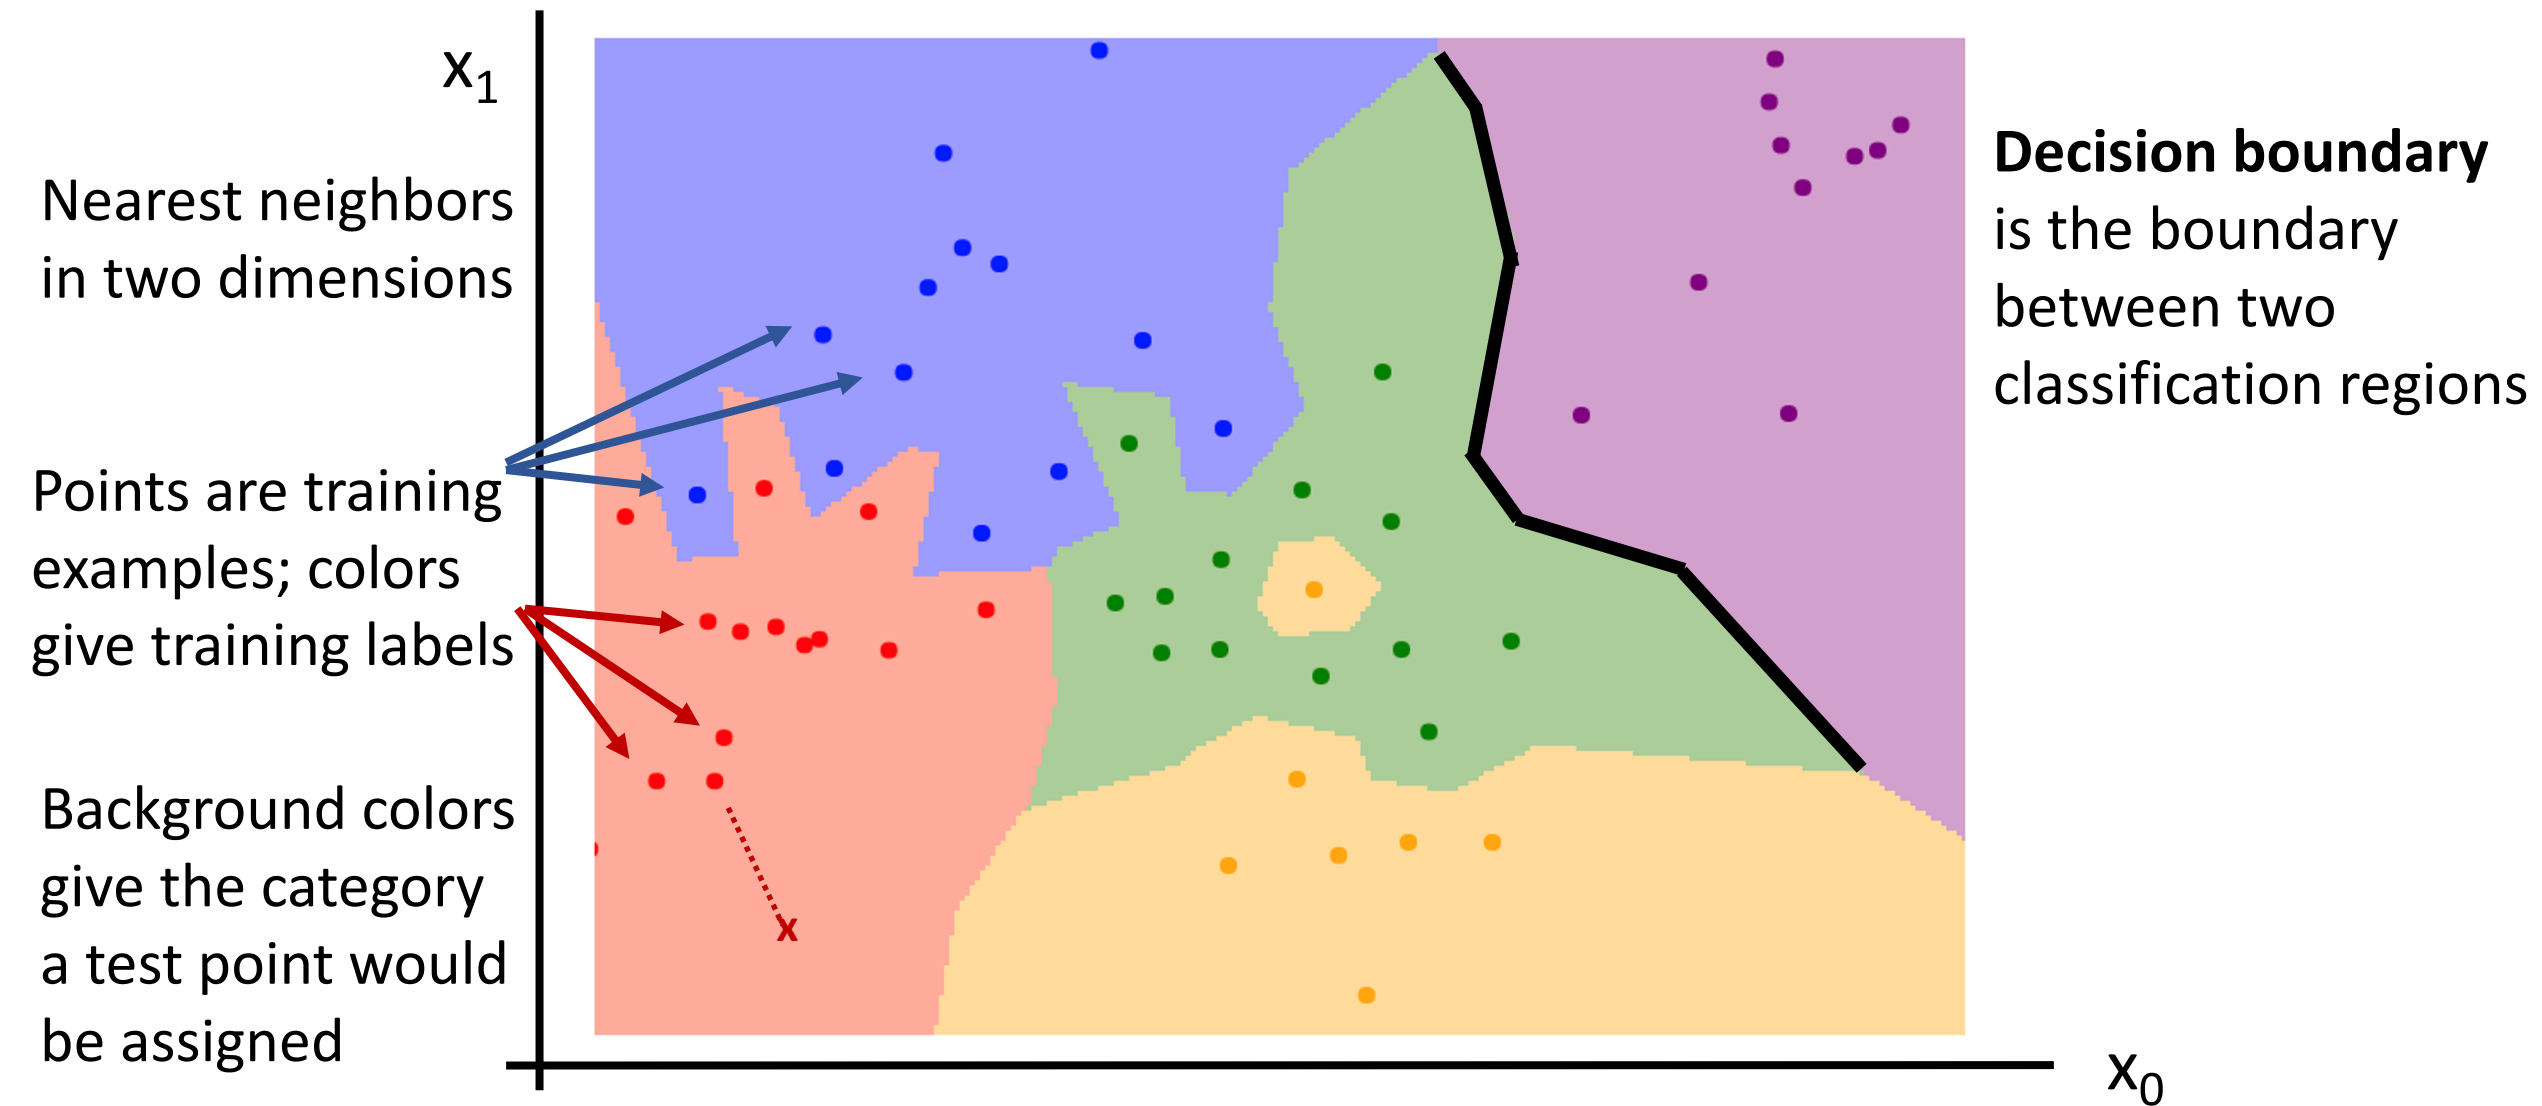
# Nearest Neighbor Decision Boundaries



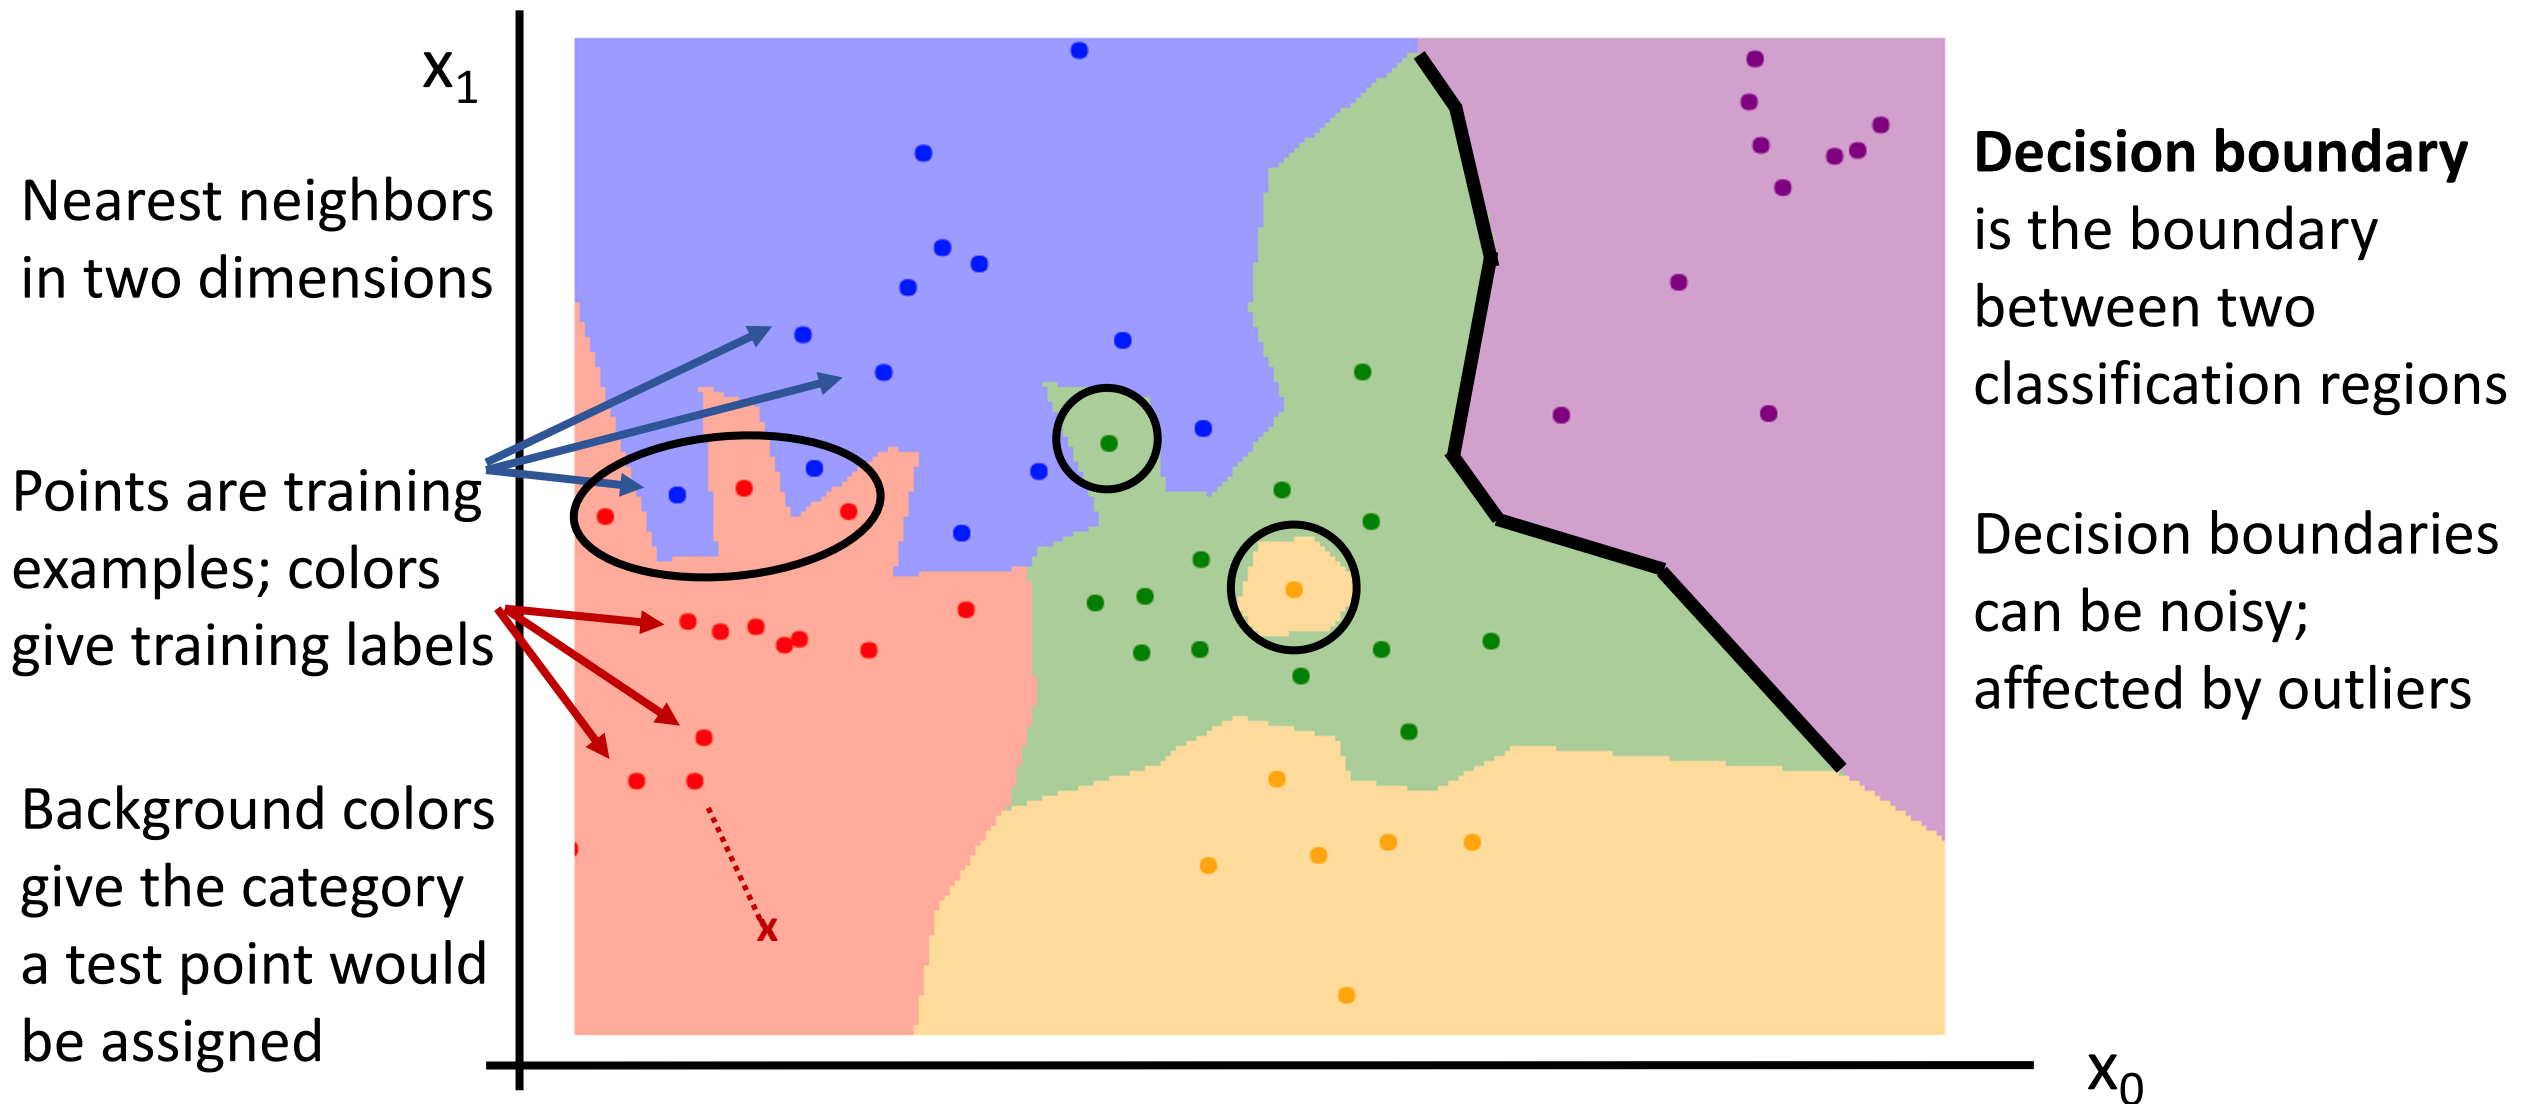
# Nearest Neighbor Decision Boundaries




# Nearest Neighbor Decision Boundaries



# Nearest Neighbor Decision Boundaries




# Nearest Neighbor Decision Boundaries




# K-Nearest Neighbors

Instead of copying label from nearest neighbor, take **majority vote** from K closest points

K = 1




K = 3




# K-Nearest Neighbors

Using more neighbors helps smooth out rough decision boundaries

$K = 1$




$K = 3$




# K-Nearest Neighbors

Using more neighbors helps reduce the effect of outliers

$K = 1$



$K = 3$







# K-Nearest Neighbors

When  $K > 1$  there can be ties between classes.  
Need to break somehow!

$K = 1$




$K = 3$



# 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


$$d_2(I_1, I_2) = \left( \sum_p (I_1^p - I_2^p)^2 \right)^{\frac{1}{2}}$$



# 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

$$d_1(I_1, I_2) = \left( \sum_p (I_1^p - I_2^p)^2 \right)^{\frac{1}{2}}$$



K = 1

# K-Nearest Neighbors: Distance Metric

Can get quite creative with the distance function, e.g. bending energy (how much do you need to transform one example to look like another)



Fig. 1. Examples of two handwritten digits. In terms of pixel-to-pixel comparisons, these two images are quite different, but to the human observer, the shapes appear to be similar.

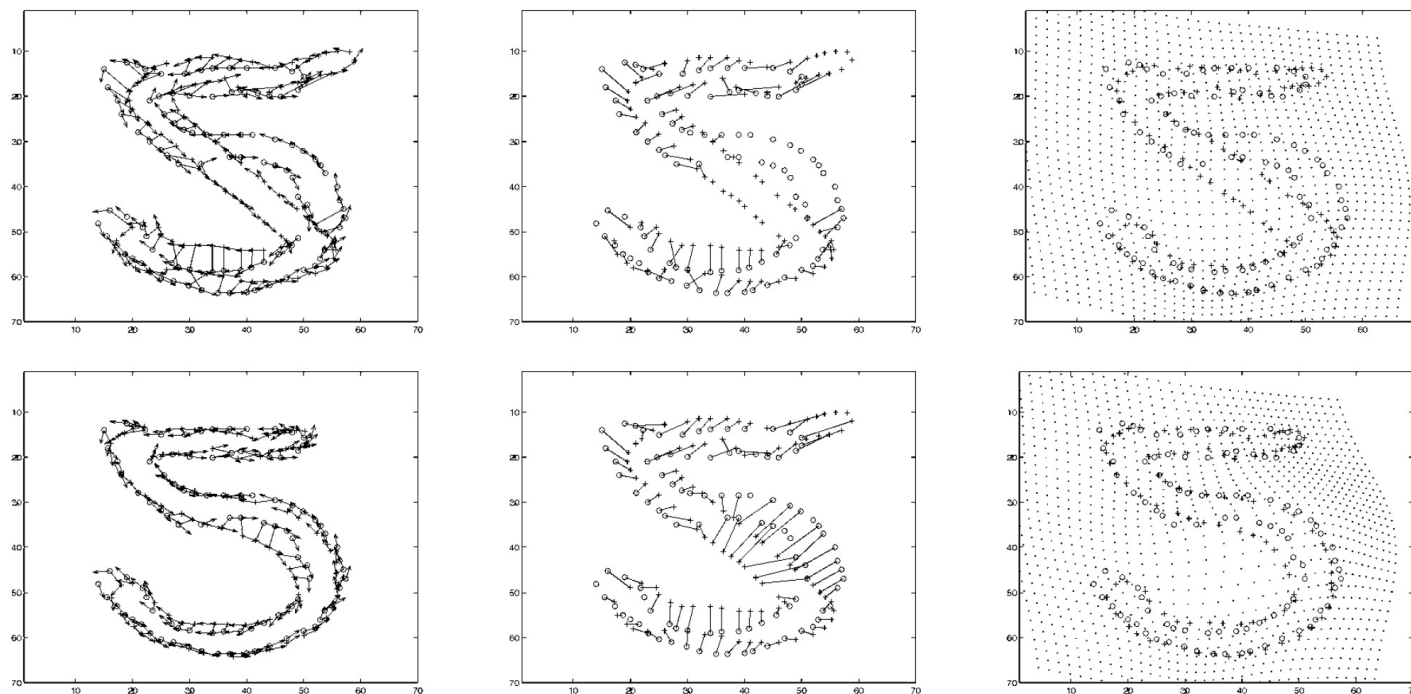


Fig. 4. Illustration of the matching process applied to the example of Fig. 1. Top row: 1st iteration. Bottom row: 5th iteration. Left column: estimated correspondences shown relative to the transformed model, with tangent vectors shown. Middle column: estimated correspondences shown relative to the untransformed model. Right column: result of transforming the model based on the current correspondences; this is the input to the next iteration. The grid points illustrate the interpolated transformation over  $\mathbb{R}^2$ . Here, we have used a regularized TPS model with  $\lambda_o = 1$ .


# K-Nearest Neighbors: Web Demo

Interactively move points around  
and see decision boundaries change

Play with L1 vs L2 metrics

Play with changing number of  
training points, value of K

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



Metric

L1 L2

Num classes

2 3 4 5

Num Neighbors (K)

1 2 3 4 5 6 7

Num points

20 30 40 50 60

# Hyperparameters

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

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

These are examples of **hyperparameters**: choices about our learning algorithm that we don't learn from the training data; instead we set them at the start of the learning process

# Hyperparameters

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

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


These are examples of **hyperparameters**: choices about our learning algorithm that we don't learn from the training data; instead we set them at the start of the learning process

Very problem-dependent. In general need to try them all and see what works best for our data / task.

# Setting Hyperparameters

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

**BAD:**  $K = 1$  always works perfectly on training data (in general, memorization is sufficient for acing the train set)






# Setting Hyperparameters


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

**BAD:**  $K = 1$  always works perfectly on training data (in general, memorization is sufficient for acing the train set)



**Idea #2:** Choose hyperparameters that work best on test data


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



# Setting Hyperparameters

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

**BAD:**  $K = 1$  always works perfectly on training data (in general, memorization is sufficient for acing the train set)



**Idea #2:** Choose hyperparameters that work best on test data

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




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

**Better!**



# 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  $k \sim 7$  works best for this data)


# K-Nearest Neighbor: Universal Approximation

As the number of training samples goes to infinity, nearest neighbor can represent any<sup>(\*)</sup> function!

(\*) Subject to many technical conditions. Only continuous functions on a compact domain; need to make assumptions about spacing of training points; etc.

# K-Nearest Neighbor: Universal Approximation


As the number of training samples goes to infinity, nearest neighbor can represent any<sup>(\*)</sup> function!



(\*) Subject to many technical conditions. Only continuous functions on a compact domain; need to make assumptions about spacing of training points; etc.

# K-Nearest Neighbor: Universal Approximation


As the number of training samples goes to infinity, nearest neighbor can represent any<sup>(\*)</sup> function!



(\*) Subject to many technical conditions. Only continuous functions on a compact domain; need to make assumptions about spacing of training points; etc.

# K-Nearest Neighbor: Universal Approximation


As the number of training samples goes to infinity, nearest neighbor can represent any<sup>(\*)</sup> function!



(\*) Subject to many technical conditions. Only continuous functions on a compact domain; need to make assumptions about spacing of training points; etc.

# K-Nearest Neighbor: Universal Approximation

As the number of training samples goes to infinity, nearest neighbor can represent any<sup>(\*)</sup> function!



(\*) Subject to many technical conditions. Only continuous functions on a compact domain; need to make assumptions about spacing of training points; etc.



# K-Nearest Neighbor on raw pixels is seldom used

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

Original



Boxed



Shifted



Tinted



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

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

# Nearest Neighbor with ConvNet features works well!



Devlin et al, "Exploring Nearest Neighbor Approaches for Image Captioning", 2015

Slide from Justin Johnson

# Can transfer more than just label!

## Image Captioning with Nearest Neighbor



A bedroom with a bed and a couch.



A cat sitting in a bathroom sink.



A train is stopped at a train station.



A wooden bench in front of a building.

Devlin et al, "[Exploring Nearest Neighbor Approaches for Image Captioning](#)", 2015

Slide from Justin Johnson


# Can transfer more than just label!

## Image Captioning with Nearest Neighbor

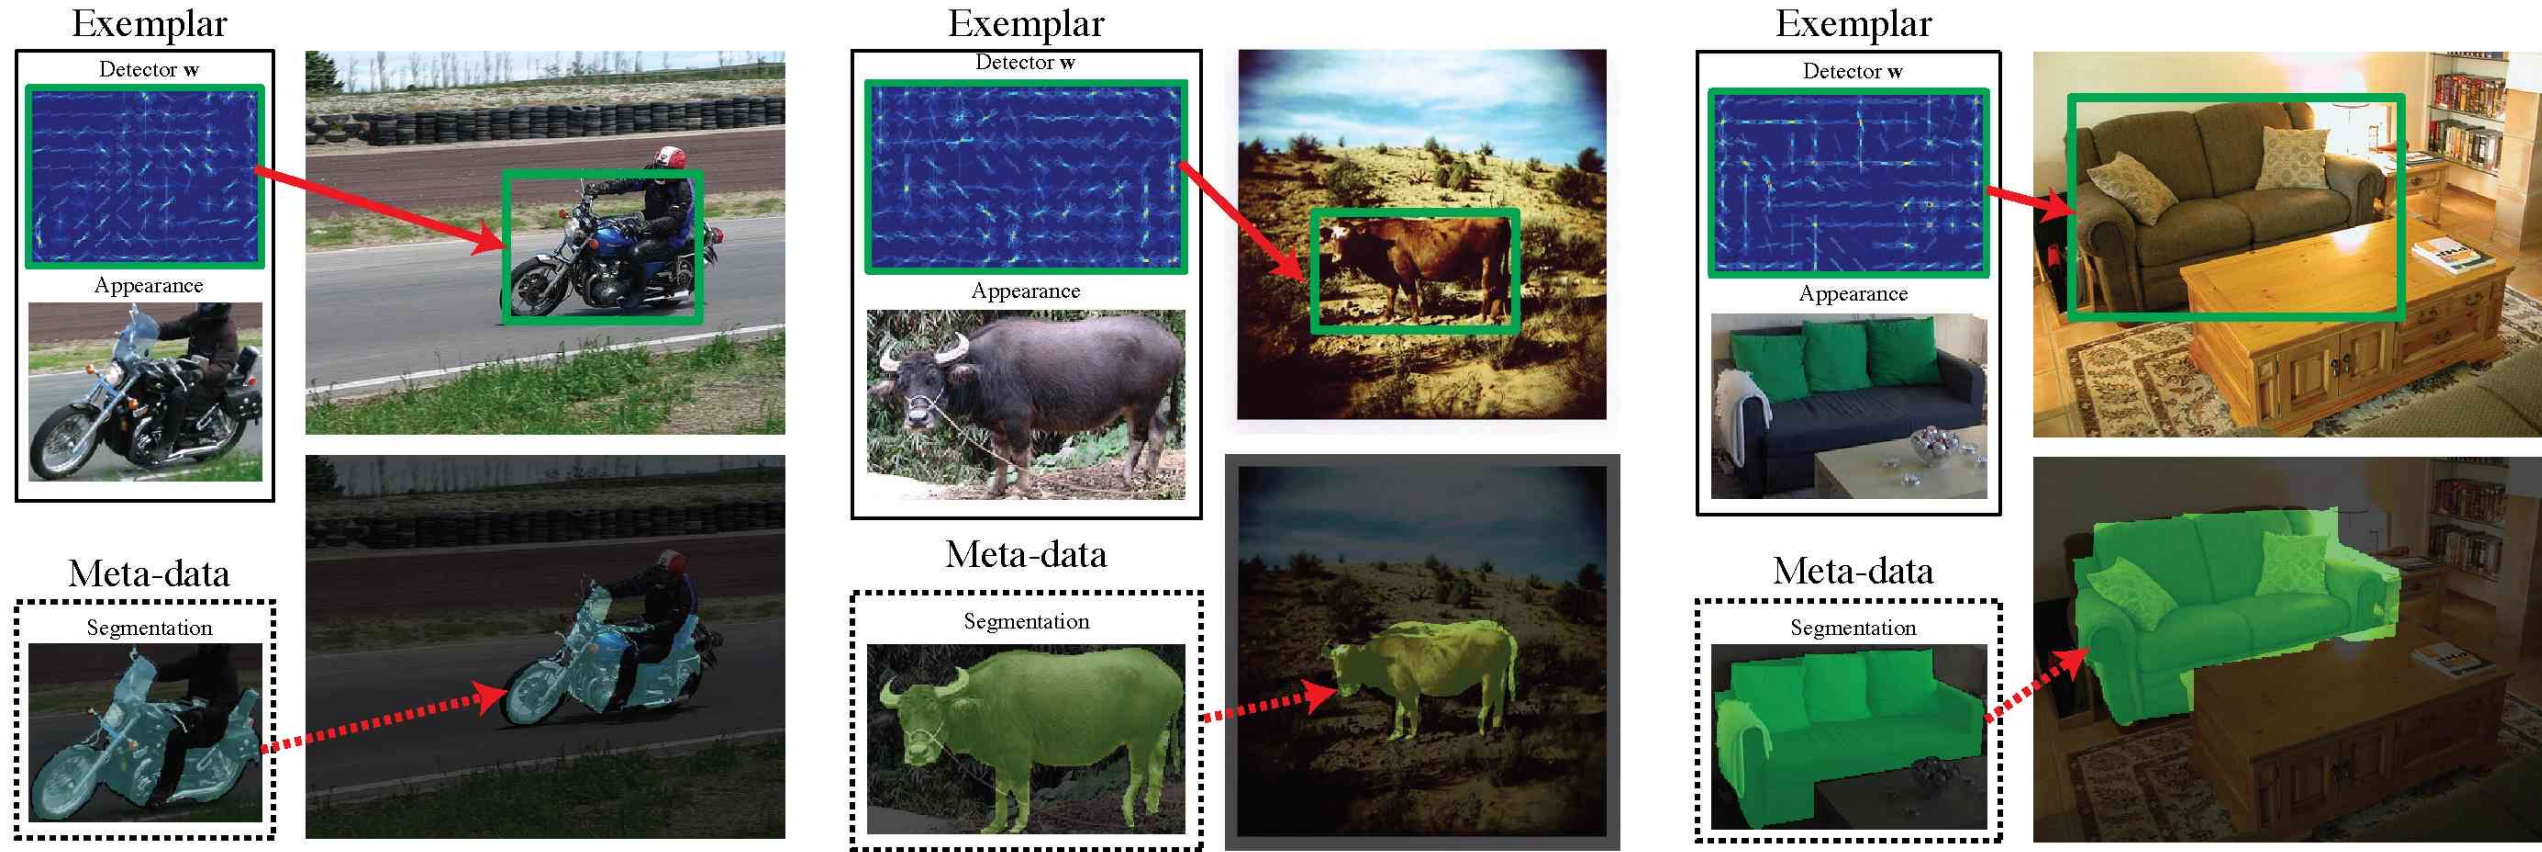
Method	BLEU 4
ME + DMSM [8]	<b>56.7</b>
LRCN [6]	53.4
Vinyals et al. [35]	53.8
Xu et al. [36]	52.3
m-RNN [25]	54.3
MLBL [18], [19]	51.7
NeuralTalk [16]	44.6
fc7-fine (CIDEr)	54.2 (2)
Human	47.1

Outperformed many other End-to-end models at the time.

# Can transfer more than just label!



# Can transfer more than just label!



# Lecture overview

- Different problems in computer vision
- Supervised classification
- Beyond supervised classification: A taxonomy of prediction problems and types of supervision
- Image classification
- k-Nearest Neighbors