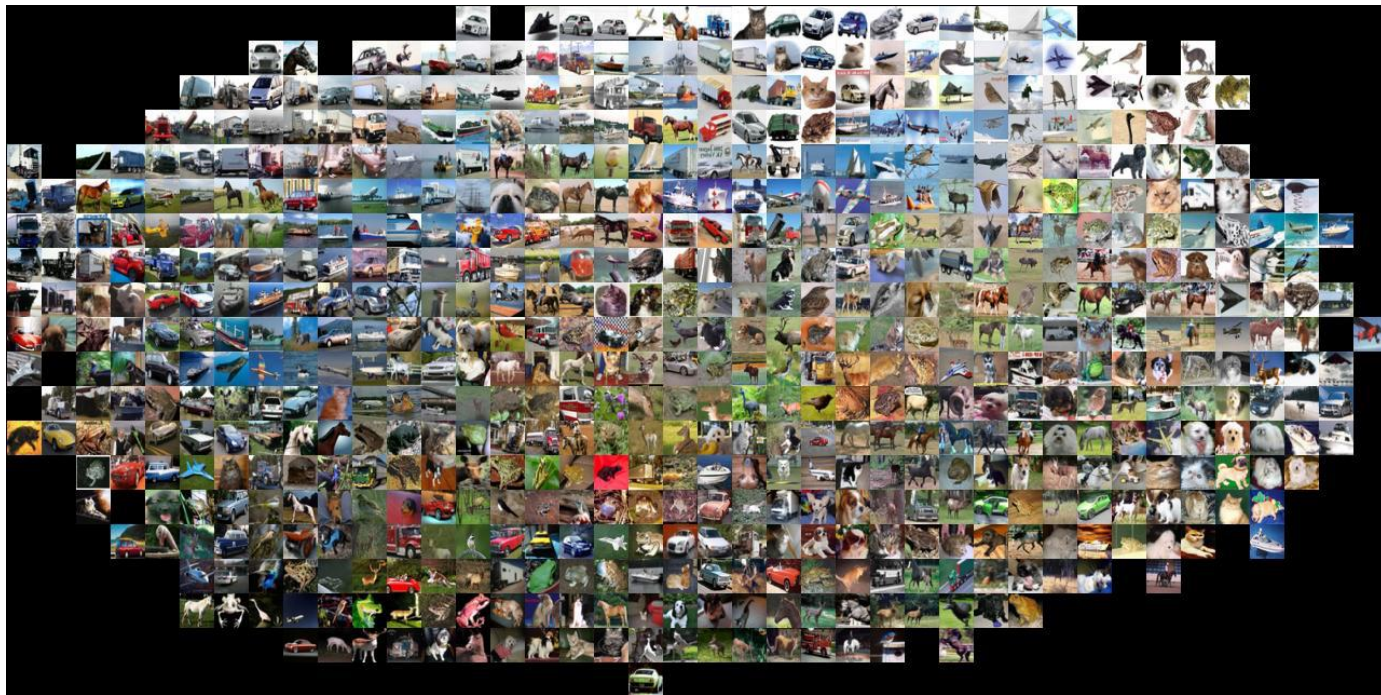


CS5670: Computer Vision

Noah Snavely

Image Classification



Announcements

- Project 4 due today
- Project 5 to be released soon
- Quiz 4 this Wednesday (4/24) (1st 10 minutes)
 - Will cover material since last quiz (photometric stereo, multi-view stereo, structure from motion, image classification (today's lecture))
- Guest lecture next Monday, 4/29, Jin Sun
 - Generative Adversarial Networks (GANs)
- Wednesday's lecture: Convolutional Neural Networks

Today

- Image classification pipeline
- Training, validation, testing
- Nearest neighbor classification
- Linear classification

- Building up to CNNs for learning
 - Next four lectures on deep learning

References

- Stanford CS231N
 - <http://cs231n.stanford.edu/>

What matters in recognition?

- Machine learning methods (e.g., linear classification, deep learning)
- Representation (e.g., SIFT, HoG, deep learned features)
- **Data** (e.g., PASCAL, ImageNet, COCO)

Image Classification:

A core task in Computer Vision

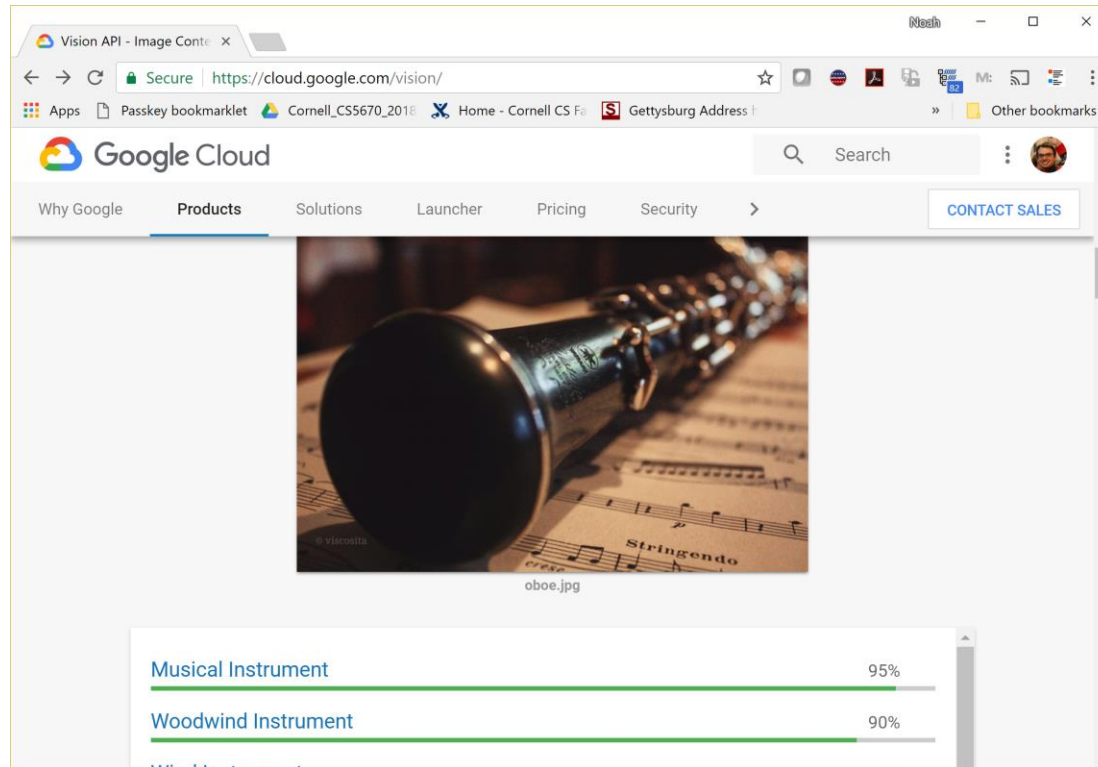
- Assume given set of discrete labels, e.g. {cat, dog, cow, apple, tomato, truck, ... }

$f(\text{apple image}) = \text{"apple"}$

$f(\text{tomato image}) = \text{"tomato"}$

$f(\text{cow image}) = \text{"cow"}$

Image classification demo



<https://cloud.google.com/vision/>

See also:

<https://aws.amazon.com/rekognition/>

<https://www.clarifai.com/>

<https://azure.microsoft.com/en-us/services/cognitive-services/computer-vision/>

Image Classification



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



cat

Image Classification: Problem



08	02	22	97	38	15	00	40	00	75	04	05	07	78	52	12	50	77	81	28
49	49	99	40	17	81	18	57	60	87	17	40	98	43	69	48	39	56	62	00
81	49	31	73	55	79	14	29	93	71	40	67	55	88	30	03	49	13	36	65
52	70	95	23	04	60	11	42	69	24	68	56	01	32	56	71	37	02	36	91
22	31	16	71	51	62	03	89	41	92	36	54	22	40	40	28	66	33	13	80
24	47	13	60	99	03	45	02	44	75	33	53	78	36	84	20	35	17	12	50
32	98	81	28	64	23	67	10	26	38	40	67	59	54	70	66	18	38	64	70
67	26	20	68	02	62	12	20	95	63	94	39	63	08	40	91	66	49	94	21
24	55	58	05	66	73	99	26	97	17	78	78	96	83	14	88	34	69	63	72
21	36	23	09	75	00	76	44	20	43	35	14	00	61	33	97	34	31	33	95
78	17	53	28	22	75	31	67	15	94	03	80	04	62	16	14	09	53	56	92
16	39	05	42	96	35	31	47	55	58	88	24	00	17	54	24	36	29	85	57
86	56	00	48	35	71	89	07	05	44	44	37	44	60	21	58	51	54	17	58
19	80	81	68	05	94	47	69	28	73	92	13	86	52	17	77	04	89	85	40
04	52	08	83	97	35	99	16	07	97	57	32	16	26	26	79	33	27	98	66
55	36	68	87	57	62	20	72	03	46	33	67	46	55	12	32	63	93	53	69
04	42	16	73	38	85	39	11	24	94	72	18	08	46	29	32	40	62	76	36
20	69	36	41	72	30	23	88	34	88	89	69	82	67	59	85	74	04	36	16
20	73	35	29	78	31	90	01	74	31	49	71	48	86	81	16	23	57	05	54
01	70	54	71	83	51	54	69	16	92	33	48	61	43	52	01	89	27	63	48

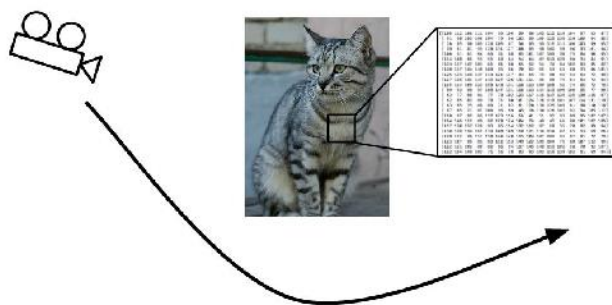
What the computer sees

image classification

82% cat
15% dog
2% hat
1% mug

Recall from last time: Challenges of recognition

Viewpoint



Illumination



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

Deformation



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

Occlusion



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

Clutter



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

Intraclass Variation



This image is [CC0 1.0](#) 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.

Data-driven approach

- Collect a database of images with labels
- Use ML to train an image classifier
- Evaluate the classifier on test images

Example training set




Data-driven approach

- Collect a database of images with labels
- Use ML to train an image classifier
- Evaluate the classifier on test images

```
def train(train_images, train_labels):  
    # build a model of images -> labels  
  
def predict(image):  
    # evaluate the model on the image  
    return class_label
```

Training

Training
Images



Training

Training
Images



Image
Features



Training
Labels



Training



Learned
Classifier

Testing

Test Image



Image
Features



Learned
Classifier



Prediction

Classifiers

- Nearest Neighbor
- kNN (“k-Nearest Neighbors”)
- Linear Classifier
- Neural Network
- Deep Neural Network
- ...

First: Nearest Neighbor (NN) Classifier

- Train
 - Remember all training images and their labels
- Predict
 - Find the closest (most similar) training image
 - Predict its label as the true label

CIFAR-10 and NN results

Example dataset: **CIFAR-10**

10 labels

50,000 training images, each image is tiny: 32x32

10,000 test images.

airplane



automobile



bird



cat



deer



dog



frog



horse



ship



truck




CIFAR-10 and NN results

Example dataset: **CIFAR-10**

10 labels

50,000 training images

10,000 test images.




For every test image (first column),
examples of nearest neighbors in rows




k-nearest neighbor

- Find the k closest points from training data
- Take **majority vote** from K closest points


the data



NN classifier



5-NN classifier



What does this look like?



What does this look like?



How to find the most similar training image? What is the distance metric?

L1 distance:

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

Where I_1 denotes image 1,
and p denotes each pixel

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

→ 456

Choice of distance metric

- Hyperparameter

L1 (Manhattan) distance

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


L2 (Euclidean) distance

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

K-Nearest Neighbors: Distance Metric

L1 (Manhattan) distance


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



K = 1

L2 (Euclidean) distance

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



K = 1


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

Hyperparameters

- What is the **best distance** to use?
- What is the **best value of k** to use?
- These are **hyperparameters**: choices about the algorithm that we set rather than learn
- How do we set them?
 - One option: try them all and see what works best

Setting Hyperparameters

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



Your Dataset

Setting Hyperparameters

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

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



Your Dataset

Setting 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



Setting 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



Setting 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

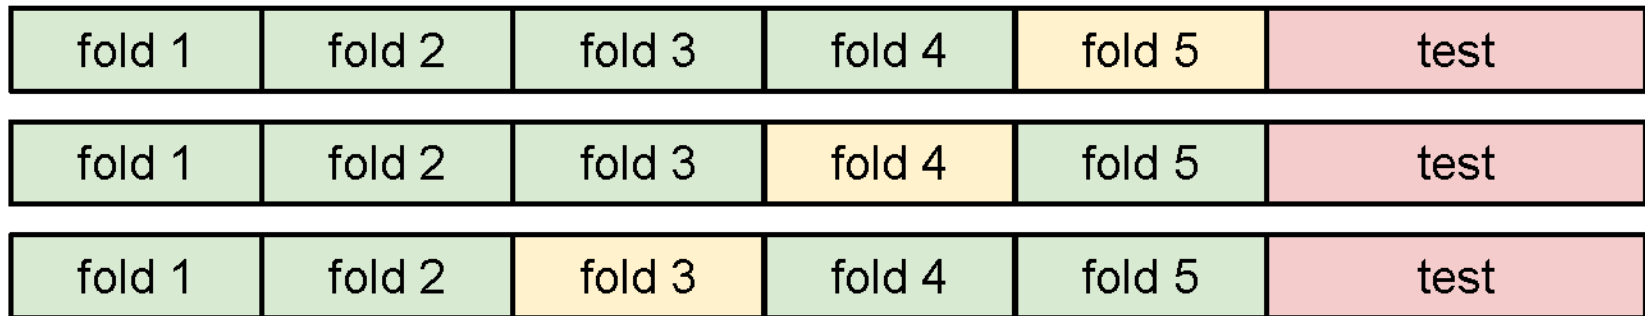
Better!




Setting Hyperparameters



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



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)

Recap: How to pick hyperparameters?


- Methodology
 - Train and test
 - Train, validate, test
- Train for original model
- Validate to find hyperparameters
- Test to understand generalizability

kNN -- Complexity and Storage

- N training images, M test images
- Training: $O(1)$
- Testing: $O(MN)$
- Hmm...
 - Normally need the opposite
 - Slow training (ok), fast testing (necessary)

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

- terrible performance at test time
- distance metrics on level of whole images can be very unintuitive



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


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!

Linear classifiers

Neural Network

Linear
classifiers



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

Score function



class scores

Score function: f

Parametric approach



image parameters

$f(\mathbf{x}, \mathbf{W})$




10 numbers,
indicating class
scores


[32x32x3]

array of numbers 0...1
(3072 numbers total)

Parametric approach: Linear classifier



Parametric approach: Linear classifier



Linear Classifier

define a **score function**

$$f(x_i, W, b) = Wx_i + b$$

The diagram illustrates the components of the linear classifier score function $f(x_i, W, b) = Wx_i + b$. Arrows point from descriptive labels to the corresponding parts of the equation: 'data (image)' points to x_i , 'weights' points to W , 'bias vector' points to b , and 'class scores' points to the function f . The term 'parameters' is also present, encompassing both W and b .

data (image)


class scores

“weights”

“bias vector”

“parameters”

Example with an image with 4 pixels, and 3 classes (cat/dog/ship)

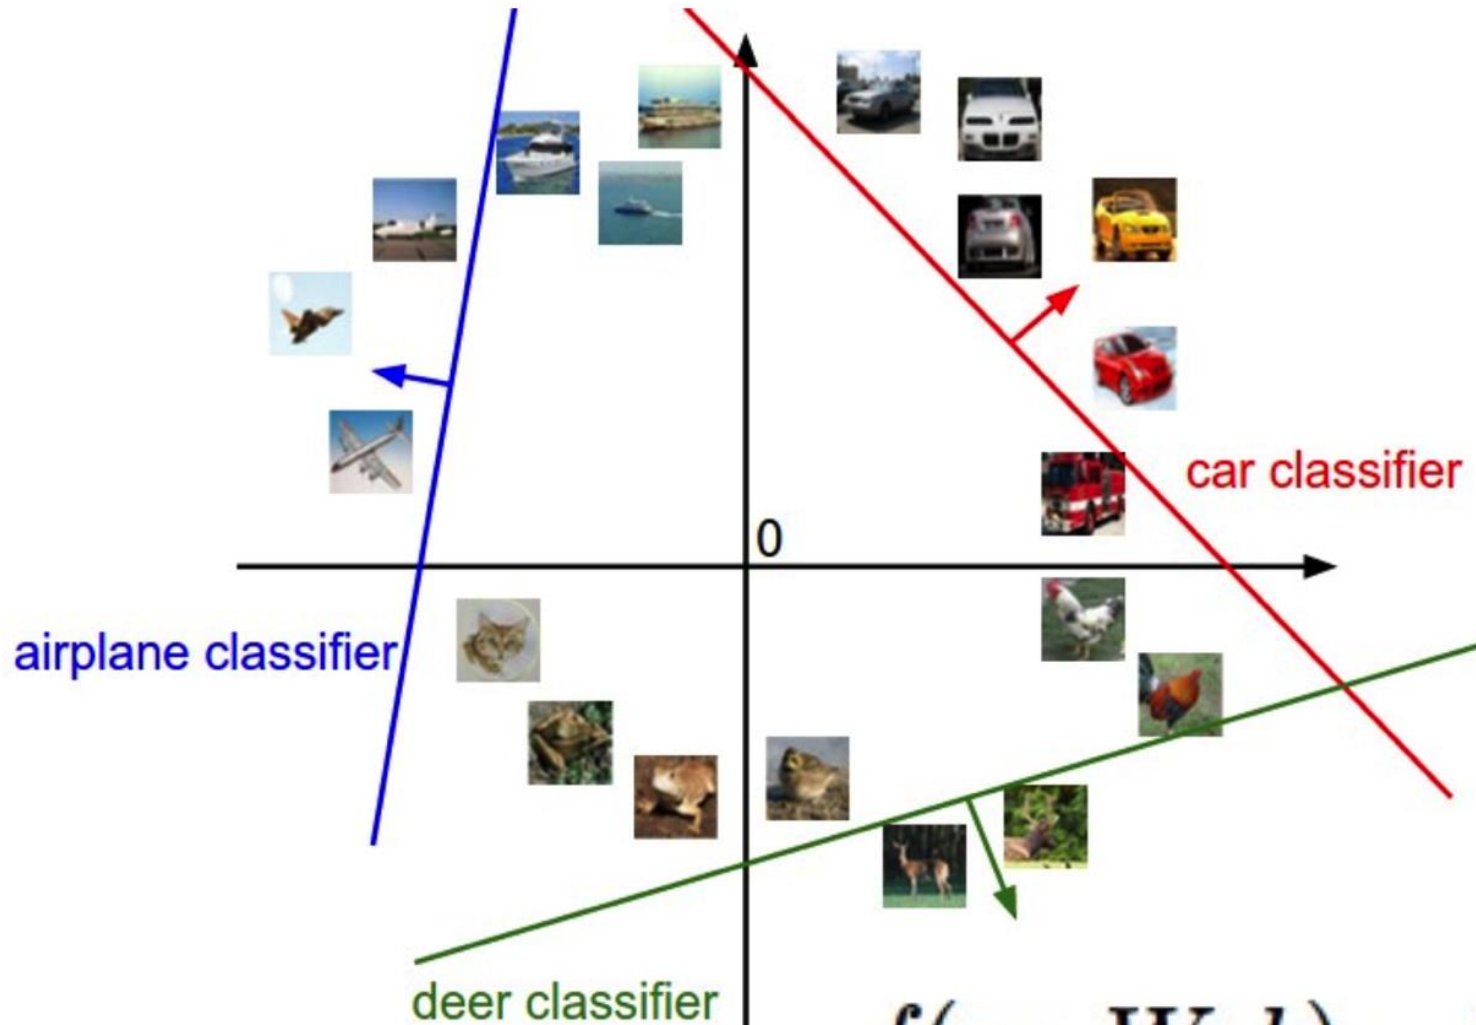


Interpretation: Template matching



$$f(x_i, W, b) = Wx_i + b$$


Geometric Interpretation



$$f(x_i, W, b) = Wx_i + b$$

Linear classifiers

- Find linear function (*hyperplane*) to separate positive and negative examples




Hard cases for a linear classifier

Class 1:

First and third quadrants

Class 2:

Second and fourth quadrants




Class 1:

$1 \leq \text{L2 norm} \leq 2$

Class 2:

Everything else




Class 1:

Three modes

Class 2:


Everything else



Linear Classifier: Three Viewpoints

Algebraic Viewpoint

$$f(x, W) = Wx$$



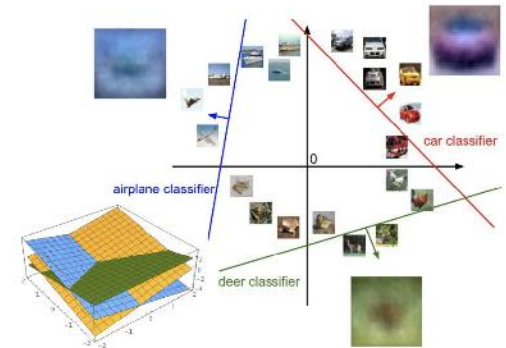
Visual Viewpoint

One template
per class



Geometric Viewpoint

Hyperplanes
cutting up space



So far: Defined a (linear) score function $f(x, W) = Wx + b$

Example class
scores for 3
images for
some W :



airplane	-3.45	-0.51	3.42
automobile	-8.87	6.04	4.64
bird	0.09	5.31	2.65
cat	2.9	-4.22	5.1
deer	4.48	-4.19	2.64
dog	8.02	3.58	5.55
frog	3.78	4.49	-4.34
horse	1.06	-4.37	-1.5
ship	-0.36	-2.09	-4.79
truck	-0.72	-2.93	6.14

[Cat image](#) by [Nikita](#) is licensed under [CC-BY 2.0](#)
[Car image](#) is [CC0 1.0](#) public domain
[Frog image](#) is in the public domain

How can we tell
whether this W
is good or bad?

Recap

- Learning methods
 - k-Nearest Neighbors
 - **Linear classification**
- Classifier outputs a **score function** giving a score to each class
- Today: how do we define how good a classifier is based on the training data? (Spoiler: define a *loss function*)

Linear classification



airplane	-3.45	-0.51	3.42
automobile	-8.87	6.04	4.64
bird	0.09	5.31	2.65
cat	2.9	-4.22	5.1
deer	4.48	-4.19	2.64
dog	8.02	3.58	5.55
frog	3.78	4.49	-4.34
horse	1.06	-4.37	-1.5
ship	-0.36	-2.09	-4.79
truck	-0.72	-2.93	6.14

TODO:

1. Define a **loss function** that quantifies our unhappiness with the scores across the training data.
2. Come up with a way of efficiently finding the parameters that minimize the loss function.
(optimization)

Output scores

[Cat image](#) by Nikita is licensed under [CC-BY 2.0](#); [Car image](#) is [CC0 1.0](#) public domain; [Frog image](#) is in the public domain

Suppose: 3 training examples, 3 classes.
With some W the scores $f(x, W) = Wx$ are:



cat	3.2	1.3	2.2
car	5.1	4.9	2.5
frog	-1.7	2.0	-3.1

A **loss function** tells how good our current classifier is

Given a dataset of examples

$$\{(x_i, y_i)\}_{i=1}^N$$

Where x_i is image and
 y_i is (integer) label

Loss over the dataset is a
sum of loss over examples:


$$L = \frac{1}{N} \sum_i L_i(f(x_i, W), y_i)$$

Loss function, cost/objective function

- Given ground truth labels (y_i), scores $f(x_i, \mathbf{W})$
 - how unhappy are we with the scores?
- Loss function or objective/cost function measures unhappiness
- During training, **want to find the parameters \mathbf{W} that minimizes the loss function**

Simpler example: binary classification

- Two classes (e.g., “cat” and “not cat”)
 - AKA “positive” and “negative” classes

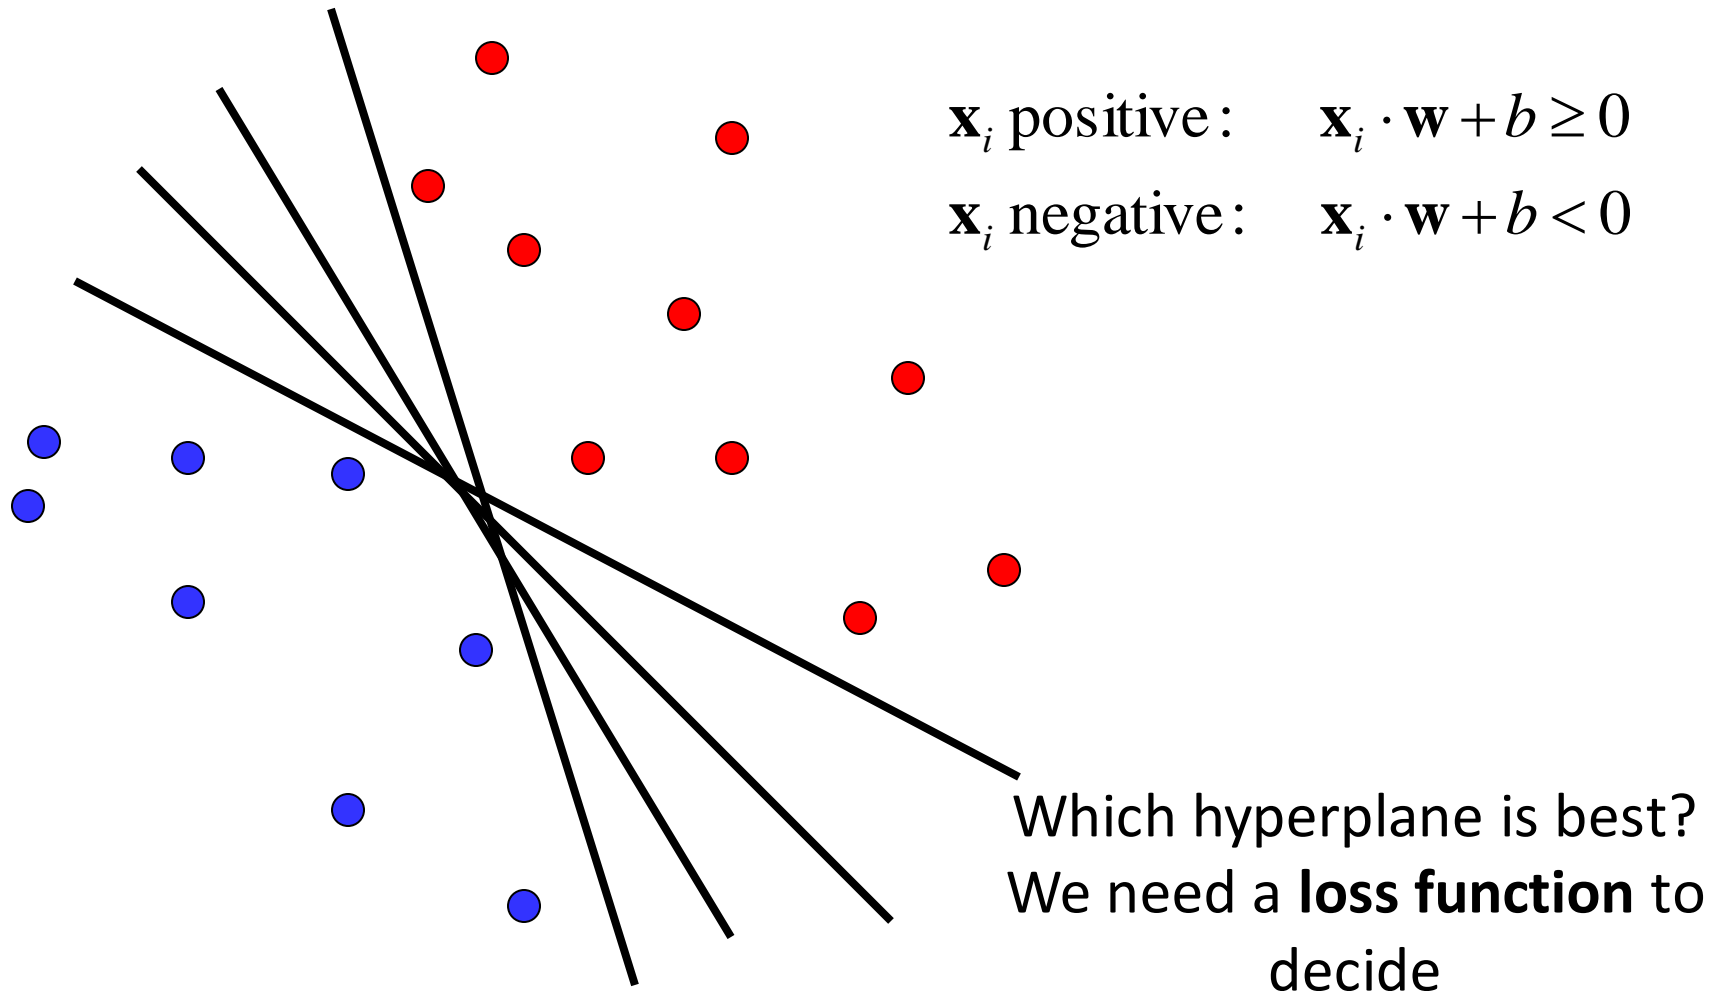


cat

not cat


Linear classifiers

- Find linear function (*hyperplane*) to separate positive and negative examples



What is a good loss function?

- One possibility
 - Number of misclassified examples



- Problems: discrete, can't break ties
- We want the loss to lead to *good generalization*
- We want the loss to work for more than 2 classes

Softmax classifier

$$f(x_i, W) = Wx_i \quad (\text{score function})$$

$$\frac{e^{f_{y_i}}}{\sum_j e^{f_j}}$$

softmax function

Example with three classes:

$$[1, -2, 0] \rightarrow [e^1, e^{-2}, e^0] = [2.71, 0.14, 1] \rightarrow [0.7, 0.04, 0.26]$$

Interpretation: squashes values into
probabilities ranging from 0 to 1

$$P(y_i \mid x_i; W)$$

Cross-entropy loss

$$f(x_i, W) = Wx_i \quad (\text{score function})$$

Losses

- Cross-entropy loss is just one possible loss function
 - One nice property is that it reinterprets scores as probabilities, which have a natural meaning
- SVM (max-margin) loss functions also used to be popular
 - But currently, cross-entropy is the most common classification loss

Summary

- Have score function and loss function
 - Currently, score function is based on linear classifier
 - Next, will generalize to convolutional neural networks
- Find W and b to minimize loss

$$L = \frac{1}{N} \sum_i -\log \left(\frac{e^{f_{y_i}}}{\sum_j e^{f_j}} \right) + \lambda \sum_k \sum_l W_{k,l}^2$$