

EECS 442 Discussion 4

Intro to Machine Learning

Discussion Agenda

- L2 Distance
- k-NN
- Train/Test Split
- Histogram of Oriented Gradients (HOG)
- Linear Classification
- Loss Functions
 - Softmax
- Gradient Descent
 - Batch Gradient Descent
 - Stochastic Gradient Descent (SGD)

Image Classification

Input Image



Output Class

dog

cat

mountain

ship

.

.

dog

cat

mountain

ship

.

.



Example dataset: CIFAR10

Images: 32x32x3 (RGB)

10 classes

Training set: 5k images per class

Test set: 1k images per class

airplane



automobile



bird



cat



deer



dog



frog



horse



ship

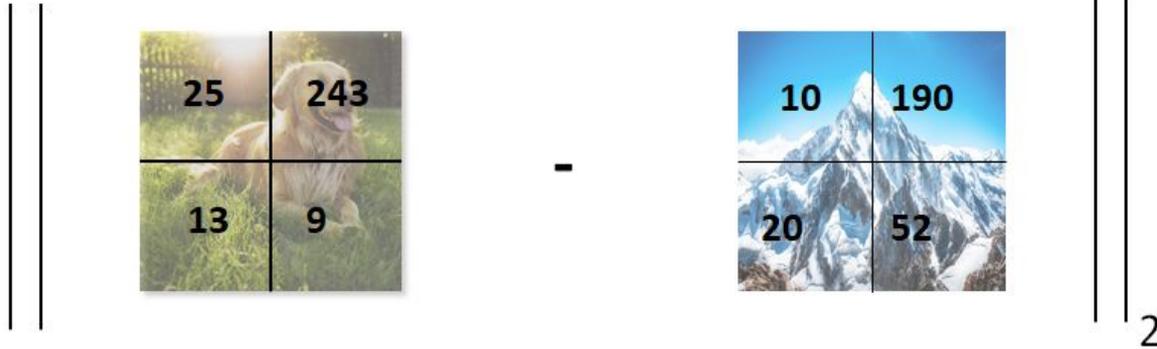


truck



L2 Distance

How do we calculate the distance between two images?



Called the L2 Distance between two images.

Square of difference

225	2809
49	43

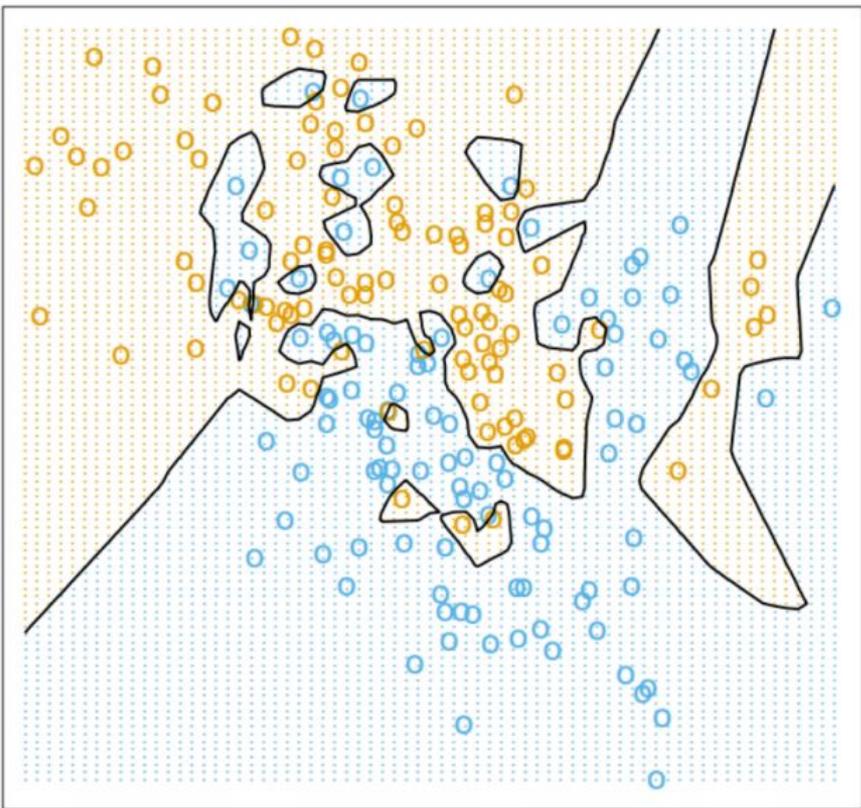
↓
Add and take square root

$$\begin{aligned} &= \sqrt{225 + 2809 + 49 + 43} \\ &= 55.91 \end{aligned}$$

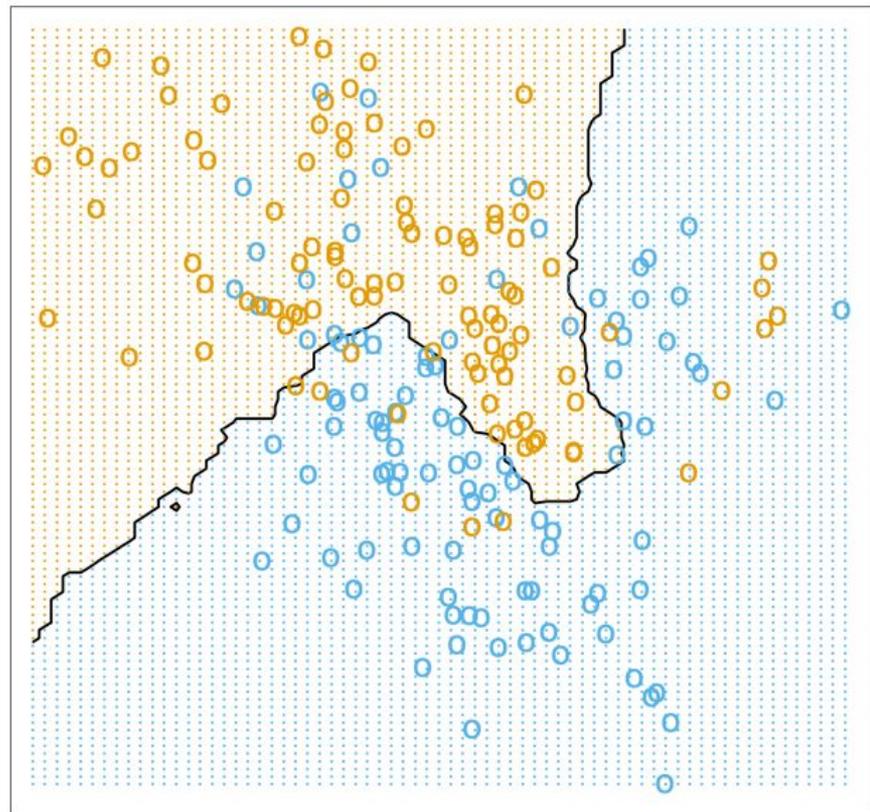
Hint for broadcasting: $\|x - y\|^2 = \|x\|^2 + \|y\|^2 - 2x^T y$

k-NN

- For each point x
- Find the k points with the shortest L2 distance from x
- Those k points are called the k -Nearest Neighbors to the point x



$k = 1$



$k = 15$

Choosing k

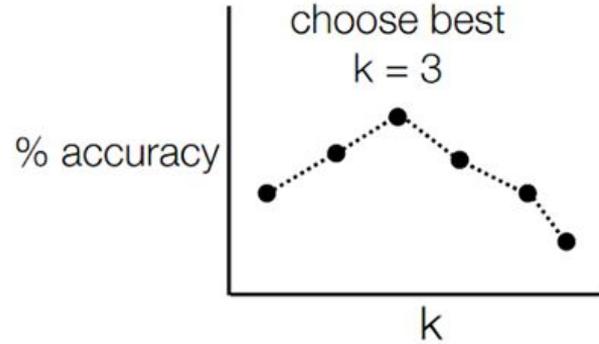
- Machine learning: generalize in the wild by training on a dataset that is representative of the samples to which we want to apply our system
- How do we get a measure of how good our system will be?
 - Calculate its accuracy on a “test set” which is done only once at the end of training
- **How do we tune our hyperparameters (like k) then?**

Train/Test Split

Pool of nearest neighbors



Training Set



Validation Set

Test Set



Choose **hyperparameters** like k , feature space, similarity function, etc.



Measures generalization.
Ideally only test once at end!

Histogram of Oriented Gradients (HOG)

- Uniquely describe the features of images
 - Edge and gradient based descriptors
1. Compute the orientations of the gradients
 2. Create a histogram of edge orientation, with votes weighted by the gradient magnitudes
 3. Perform block normalization across the histogram

Input image



Histogram of Oriented Gradients

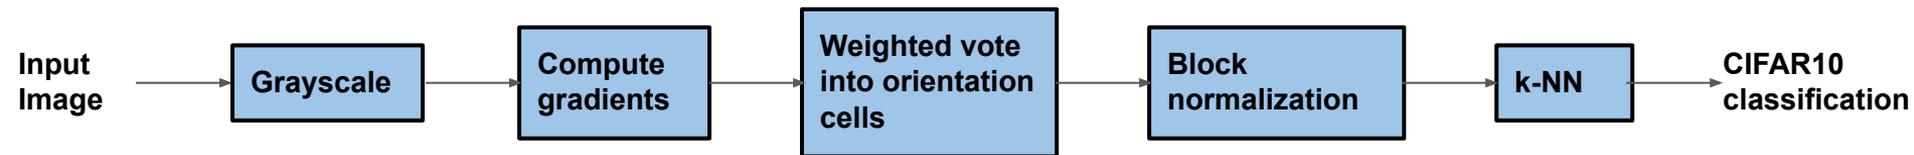


HOG Classification Workflow

Original HOG



Our simplified HOG

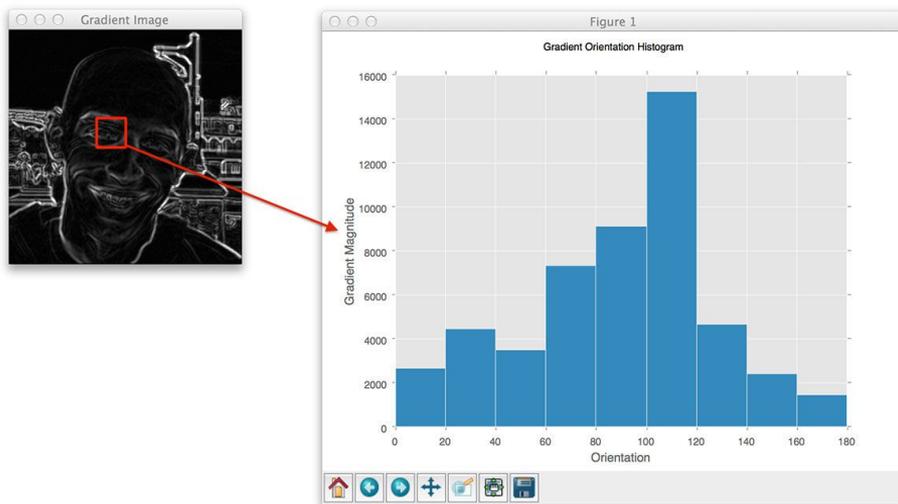


Computing Orientations of the Image Gradients

- Compute gradients of an image along the horizontal and vertical directions
 - Convolve with $dx = [1 \ -1]$ and $dy = [1 \ -1]^T$ from ps1
 - or convolve with a Sobel filter
- Compute magnitude of each gradient
 - $\text{magnitude} = \sqrt{G_x^2 + G_y^2}$
- Compute orientation/angle of each gradient
 - $\text{orientation} = \tan^{-1}(G_y, G_x)$
 - Convert angles to degrees in range $[0, 180 \text{ deg}]$
 - Why $[0, 180 \text{ deg}]$ instead of $[0, 360 \text{ deg}]$?

Create Histogram from Orientations and Magnitudes

- Iterate through each pixel in every cell
- Calculate vote for a bin based on the orientation of the gradient element centred on it (with linear interpolation) $\text{vote} = \text{magnitude} * \text{weight}$
- Weight vote based on magnitude and place into bins



Linear Interpolation

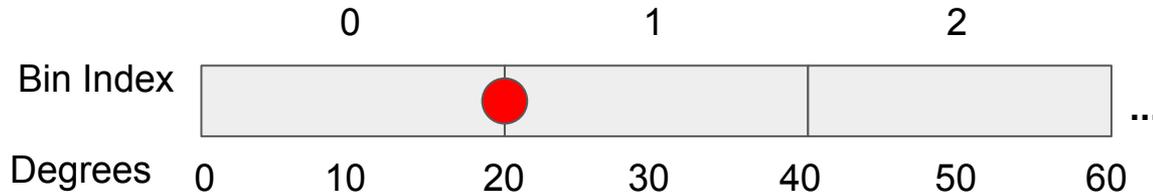
- Weigh a magnitude by how far away it is from the center of a bin
- $\text{weight} = \text{dist_from_bin_center} / \text{step_size}$

i.e. weight for 20 degrees (between bin 0 and 1)

$\text{lower_bin_index} = \text{floor}(\text{angle} / \text{step_size} - 0.5) = 0$

weight for bin 0: $(\text{angle} - (\text{lower_index} + 0.5) * \text{step_size}) / \text{step_size} = 0.5$

weight for bin 1: $((\text{lower_index} + 1.5) * \text{step_size} - \text{angle}) / \text{step_size} = 0.5$



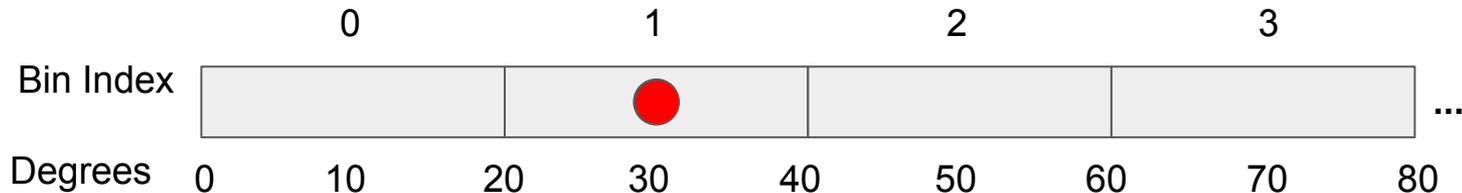
More Examples of Linear Interpolation

Weight for 30 degrees (exactly at center of bin 1)

$\text{lower_bin_index} = \text{floor}(\text{angle} / \text{step_size} - 0.5) = 1$

weight for bin 1: $(\text{angle} - (\text{lower_index} + 0.5) * \text{step_size}) / \text{step_size} = 0$

weight for bin 2: $((\text{lower_index} + 1.5) * \text{step_size} - \text{angle}) / \text{step_size} = 1$



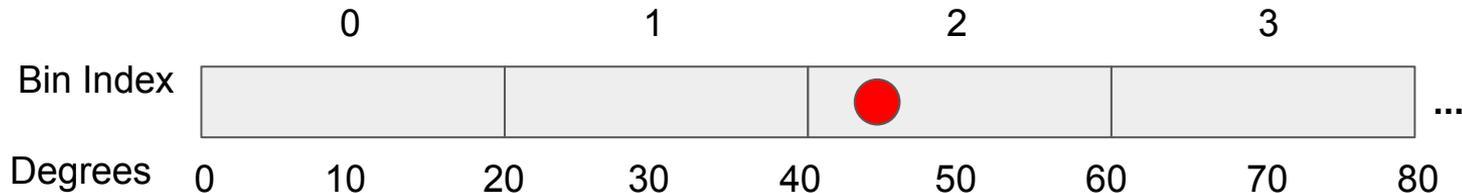
More Examples of Linear Interpolation

Weight for 45 degrees (between bins 1 and 2)

$\text{lower_bin_index} = \text{floor}(\text{angle} / \text{step_size} - 0.5) = 1$

weight for bin 1: $(\text{angle} - (\text{lower_index} + 0.5) * \text{step_size}) / \text{step_size} = 0.75$

weight for bin 2: $((\text{lower_index} + 1.5) * \text{step_size} - \text{angle}) / \text{step_size} = 0.25$

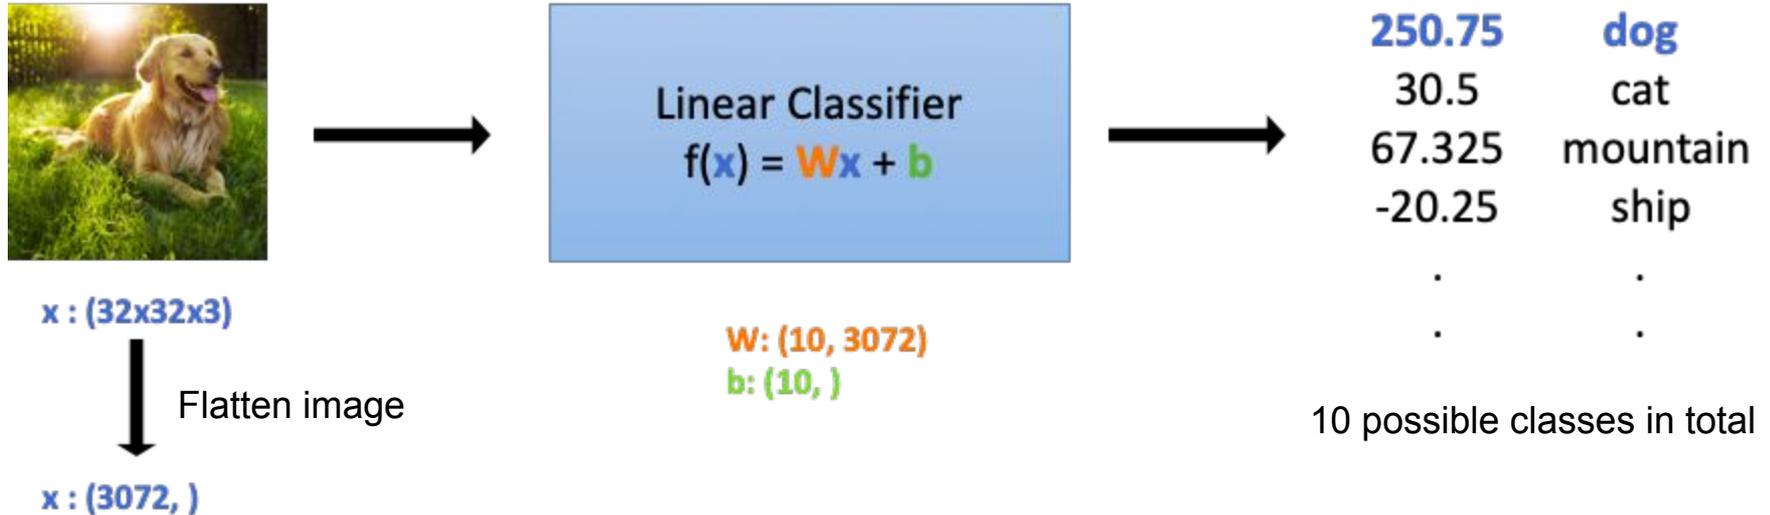


Block Normalization

- Each block consists of $n \times n$ cells
- Divide everything in a block by the square root of the sum of the squares
 - Plus an epsilon term to avoid division by zero
- Keep regions of image at the same intensity/brightness

$$L2\text{-norm, } \mathbf{v} \rightarrow \mathbf{v} / \sqrt{\|\mathbf{v}\|_2^2 + \epsilon^2}$$

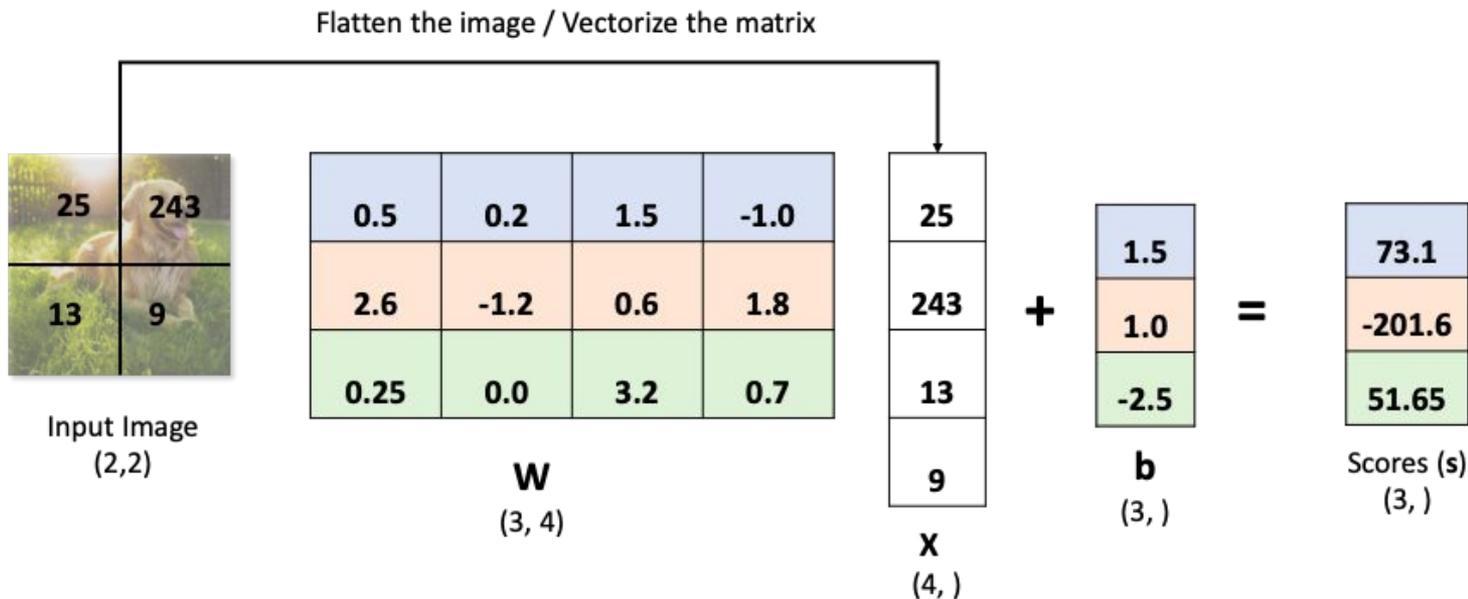
Linear classifiers



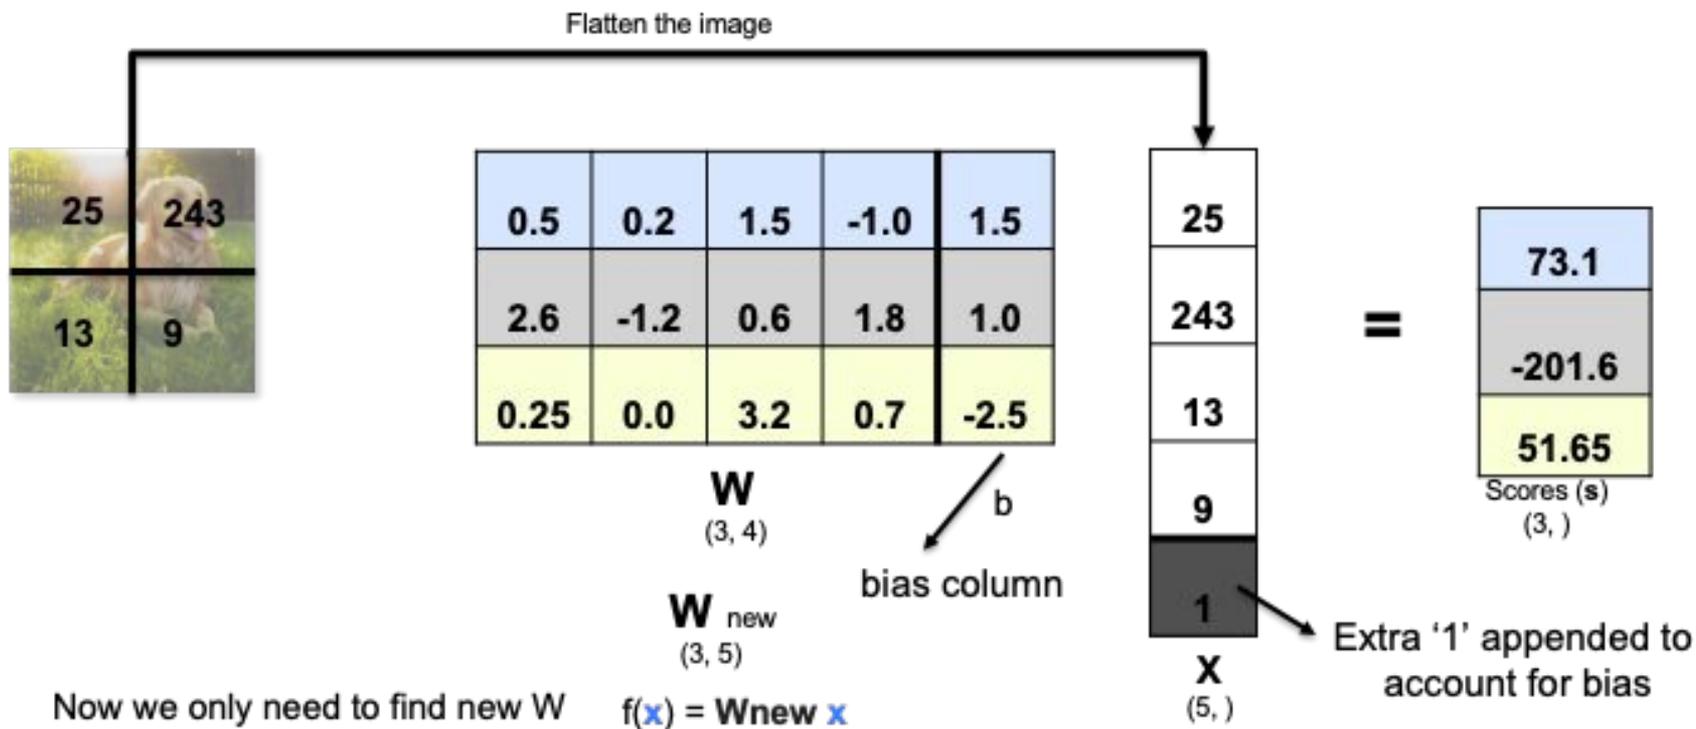
```
img = img.reshape(img.shape[0], -1)
```

Linear classifiers

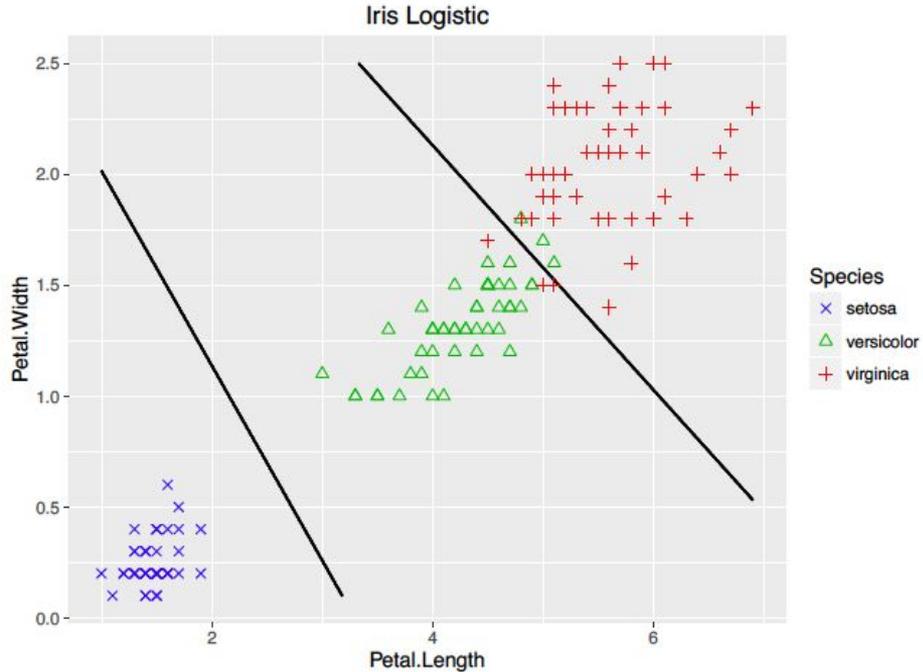
Example 2x2 image, 3 classes (dog, cat, mountain)



Bias Trick



Linear classifier boundary



Loss Function

Computes how much prediction from current model deviates from target

Cross entropy

$$\mathcal{L}(\hat{\mathbf{y}}, \mathbf{y}) = H(\mathbf{y}, \hat{\mathbf{y}}) = - \sum_{k=1}^K y_k \log \hat{y}_k \quad \leftarrow \text{easy to optimize, good approximation}$$

Softmax loss / Multinomial logistic loss

S: score, y: one of the labels



For a single image, x
with label ' y '

$$L(s, y) = -\log \frac{e^{s_{y_i} - \max(s_j)}}{\sum_{j=1}^C e^{s_j - \max(s_j)}}$$

Dog	2.5
Cat	-1.3
Mountain	5.2

Scores, $s = Wx$

exp

12.18
0.27
181.27

$$e^{s_y}$$

$y = 1, 2, 3$

1: Dog

2: Cat

3: Mountain

normalize

Probabilities

0.06
0.0
0.94

In this case, class $y = 1$ (Dog)

$$L(s, 1) = -\log 0.06 = 2.81$$

$$\frac{e^{s_y}}{\sum_{j=1}^3 e^{s_j}}$$

Softmax loss properties

$$L(s, y) = -\log \frac{e^{s_{y_i} - \max(s_j)}}{\sum_{j=1}^C e^{s_j - \max(s_j)}} \quad \text{C: Number of classes}$$

What is the min/max possible Value of $L(s, y)$? \longrightarrow Min 0, Max +infinity

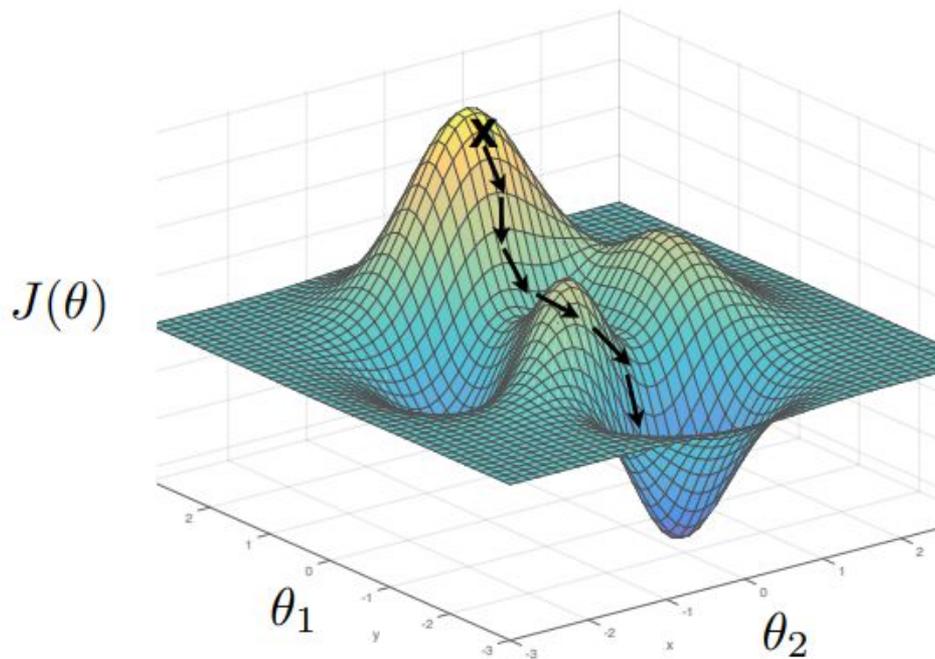
If all scores were really small and thus approximately the same, \longrightarrow $\log(C)$
What would be the value of $L(s, y)$?

Gradient Descent

- Want to optimize some objective function J

$$\theta^* = \arg \min_{\theta} \underbrace{\sum_{i=1}^N \mathcal{L}(f_{\theta}(\mathbf{x}_i), \mathbf{y}_i)}_{J(\theta)}$$

Gradient Descent



Gradient Descent

One iteration of gradient descent:

$$\theta^{t+1} = \theta^t - \eta_t \nabla_{\theta} J(\theta) \Big|_{\theta = \theta^t}$$

 learning rate

Batch Gradient Descent

Loss function:

$$J(\theta) = \frac{1}{N} \sum_{i=1}^N L(x_i, y_i, \theta)$$

Its gradient is the **sum of gradients** for each example:

$$\nabla J(\theta) = \frac{1}{N} \sum_{i=1}^N \nabla L(x_i, y_i, \theta)$$

Requires iterating over every training example each gradient step!

Can we speed this up?

Stochastic Gradient Descent

- Sample a point instead of looping over all training examples

$$\nabla J(\theta) \approx \frac{1}{|B|} \sum_{i \in B} \nabla L(x_i, y_i, \theta)$$

where B is a minibatch: a random subset of examples