

Artificial Neurons

Danna Gurari

University of Colorado Boulder
Spring 2025



<https://dannagurari.colorado.edu/course/neural-networks-and-deep-learning-spring-2025/>

Review

- Last lecture:
 - Deep learning applications
 - History of neural networks and deep learning
 - How does a machine learn?
 - Course logistics
- Please keep up with assigned readings posted to course website
- Questions?

Today's Topics


- Supervised learning: approach to develop a model
- Artificial neuron model: basic unit of neural networks
- Evaluating classification models

Today's Topics

- Supervised learning: approach to develop a model
- Artificial neuron model: basic unit of neural networks
- Evaluating classification models

Goal: Design Models that **Generalize** Well to New, Previously Unseen Examples

Example:



Label:

Hairy

Hairy

Not Hairy



Hairy



Goal: Design Models that **Generalize Well** to New, Previously Unseen Examples

1. Split data into a “**training set**” and “**test set**”

Training Data

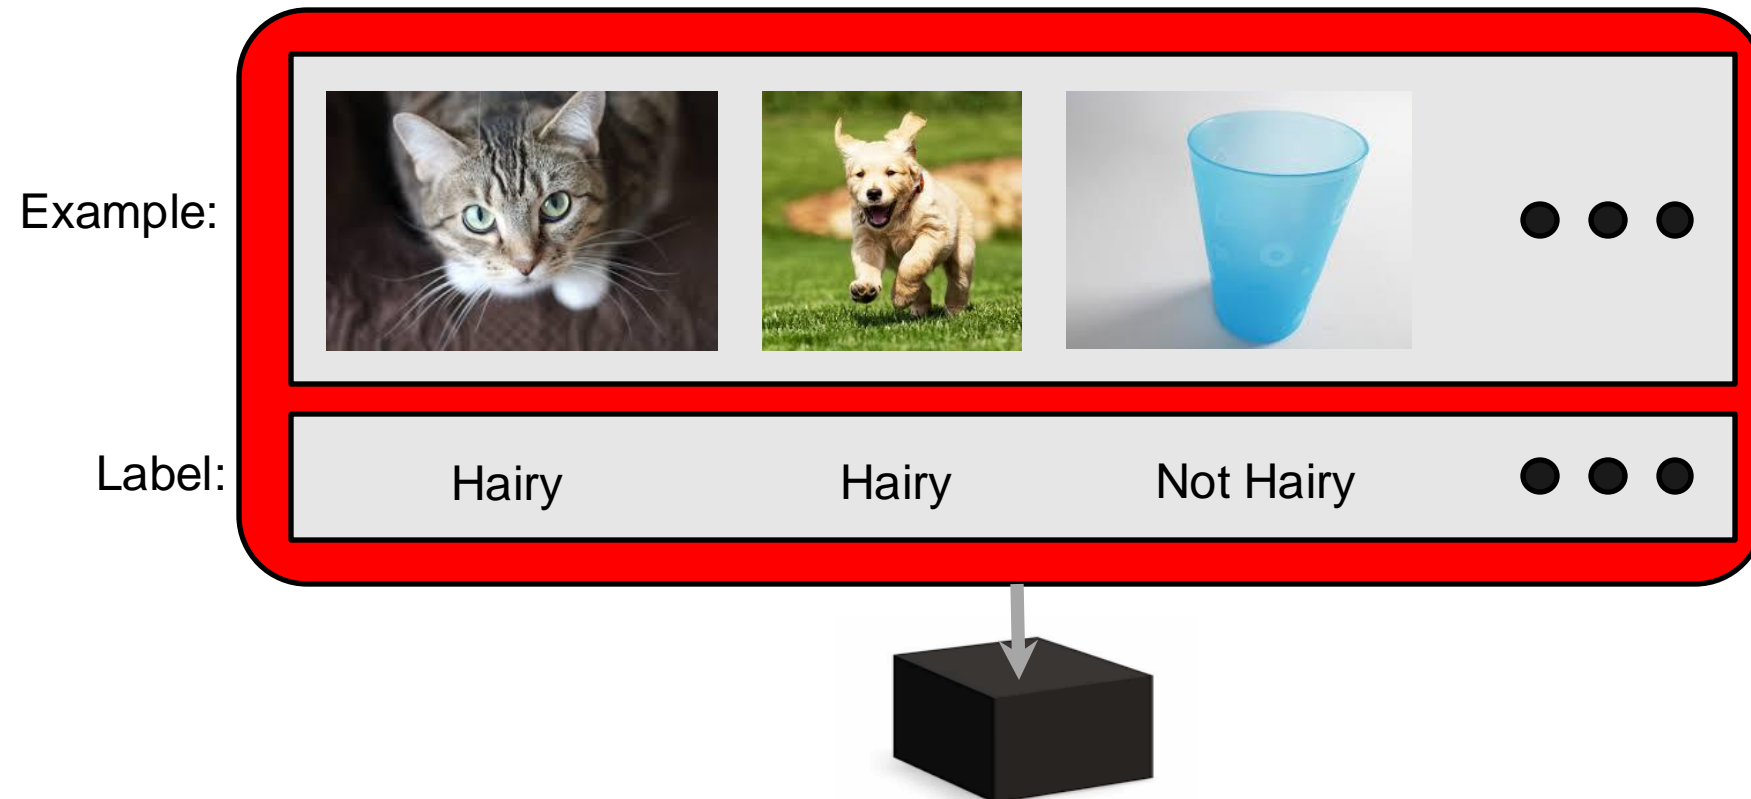
Test Data



Goal: Design Models that **Generalize Well** to New, Previously Unseen Examples


2. Train model on “**training set**” to try to minimize prediction error on it

Training Data



Goal: Design Models that **Generalize** Well to New, Previously Unseen Examples

3. Apply trained model on “**test set**” to measure generalization error



Example:



Label:

Hairy




Predicted Label:

?





Test Data



Goal: Design Models that **Generalize** Well to New, Previously Unseen Examples

3. Apply trained model on “**test set**” to measure generalization error



Example:



Label:

Hairy




Predicted Label:

Not Hairy





Test Data



Goal: Design Models that **Generalize** Well to New, Previously Unseen Examples

3. Apply trained model on “**test set**” to measure generalization error



Example:



Label:

Hairy



Predicted Label:

Not ~~X~~airy



Test Data



Supervised Learning Mimics One of the Ways that Humans Can Learn; e.g.,

Supervised Learning

(identify patterns from *structured* data with labels of target outputs)



Unsupervised Learning

(identify patterns in *unstructured* data)



Today's Topics

- Supervised learning: approach to develop a model
- **Artificial neuron model: basic unit of neural networks**
- Evaluating classification models

Vision

New York Times article, July 8, 1958 :

<https://www.nytimes.com/1958/07/08/archives/new-navy-device-learns-by-doing-psychologist-shows-embryo-of.html>

NEW NAVY DEVICE LEARNS BY DOING

Psychologist Shows Embryo
of Computer Designed to
Read and Grow Wiser

WASHINGTON, July 7 (UPI)

—The Navy revealed the embryo of an electronic computer today that it expects will be able to walk, talk, see, write, reproduce itself and be conscious of its existence.

The embryo—the Weather Bureau's \$2,000,000 "704" computer—learned to differentiate between right and left after fifty attempts in the Navy's demonstration for newsmen.

The service said it would use this principle to build the first of its Perceptron thinking machines that will be able to read

and write. It is expected to be finished in about a year at a cost of \$100,000.

Idea: Mimic Human Machinery

Neuron: basic computing machinery that enables human behavior!

- receives, processes, and transmits information, including:

“hot”



<https://www.clipart.email/clipart/dont-touch-hot-stove-clipart-73647.html>

“loud”




<https://kisselpaso.com/if-the-sun-city-music-fest-gets-too-loud-there-is-a-phone-number-you-can-call-to-complain/>

“spicy”



https://www.babycenter.com/404_when-can-my-baby-eat-spicy-foods_1368539.bc

Idea: Mimic Human Machinery



<https://www.youtube.com/watch?v=oa6rvUJlg7o>

Idea: Mimic Human Machinery



Sidenote: It Remains An Open Research Problem to Understand How Individual Neurons Work



- When the input signals exceed a certain threshold within a short period of time, a neuron “fires”
- Neuron “firing” is an “all-or-none” process, where either a signal is sent or nothing happens

Historical Context: Artificial Neurons

First mathematical
model of neuron

1943

1945

1950

1956

1959



First programmable
machine

Turing test

AI


Machine
learning



Warren McCulloch
(Neurophysiologist)

http://web.csulb.edu/~cwallis/artificialn/warren_mcculloch.html


Emerged from
interdisciplinary
collaboration



Walter Pitts
(Mathematician)

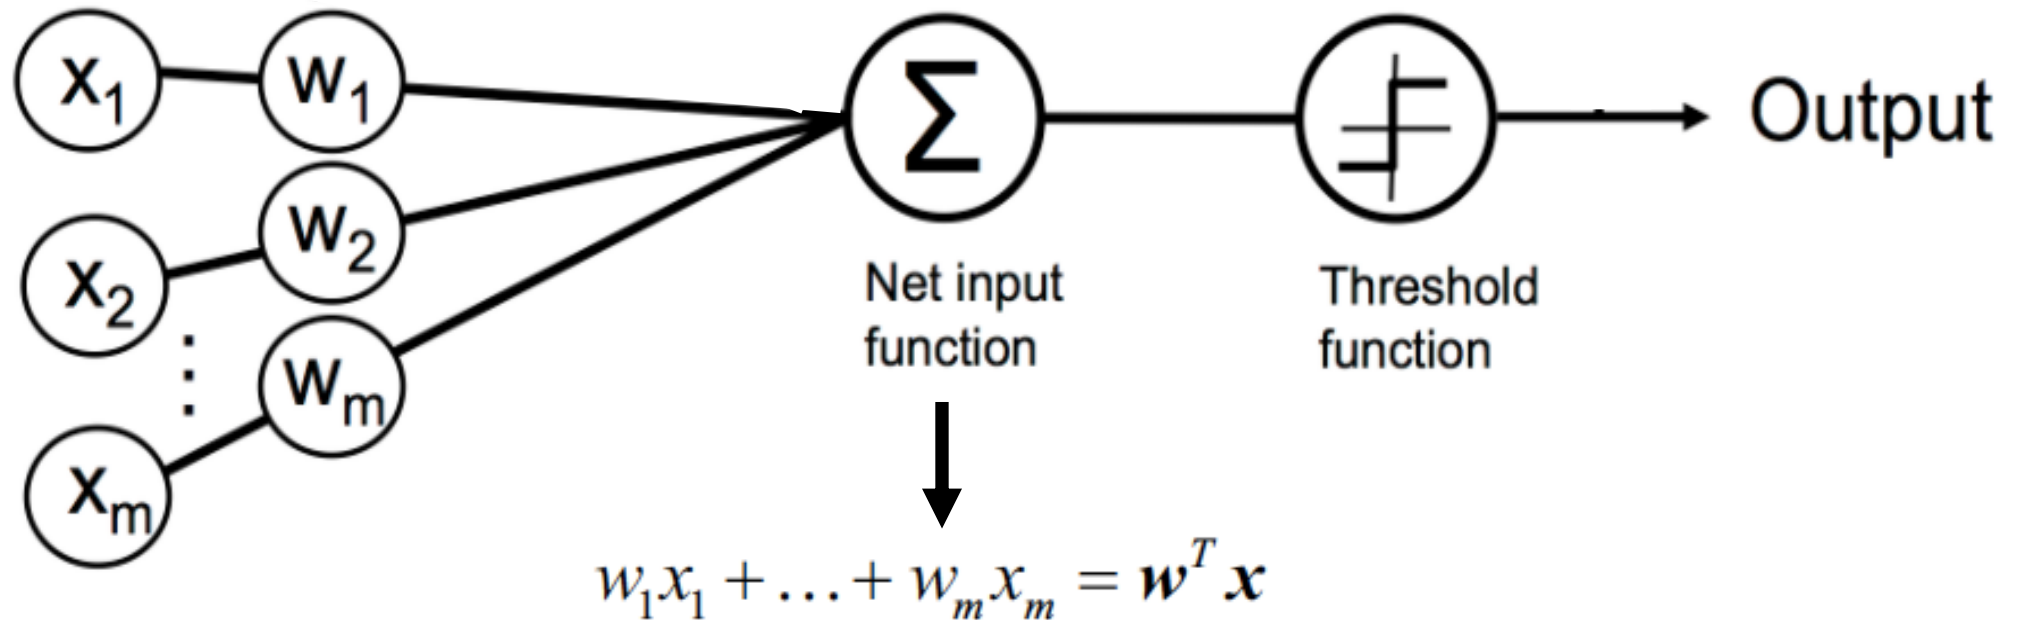
https://en.wikipedia.org/wiki/Walter_Pitts

Artificial Neuron: McCulloch-Pitts Neuron



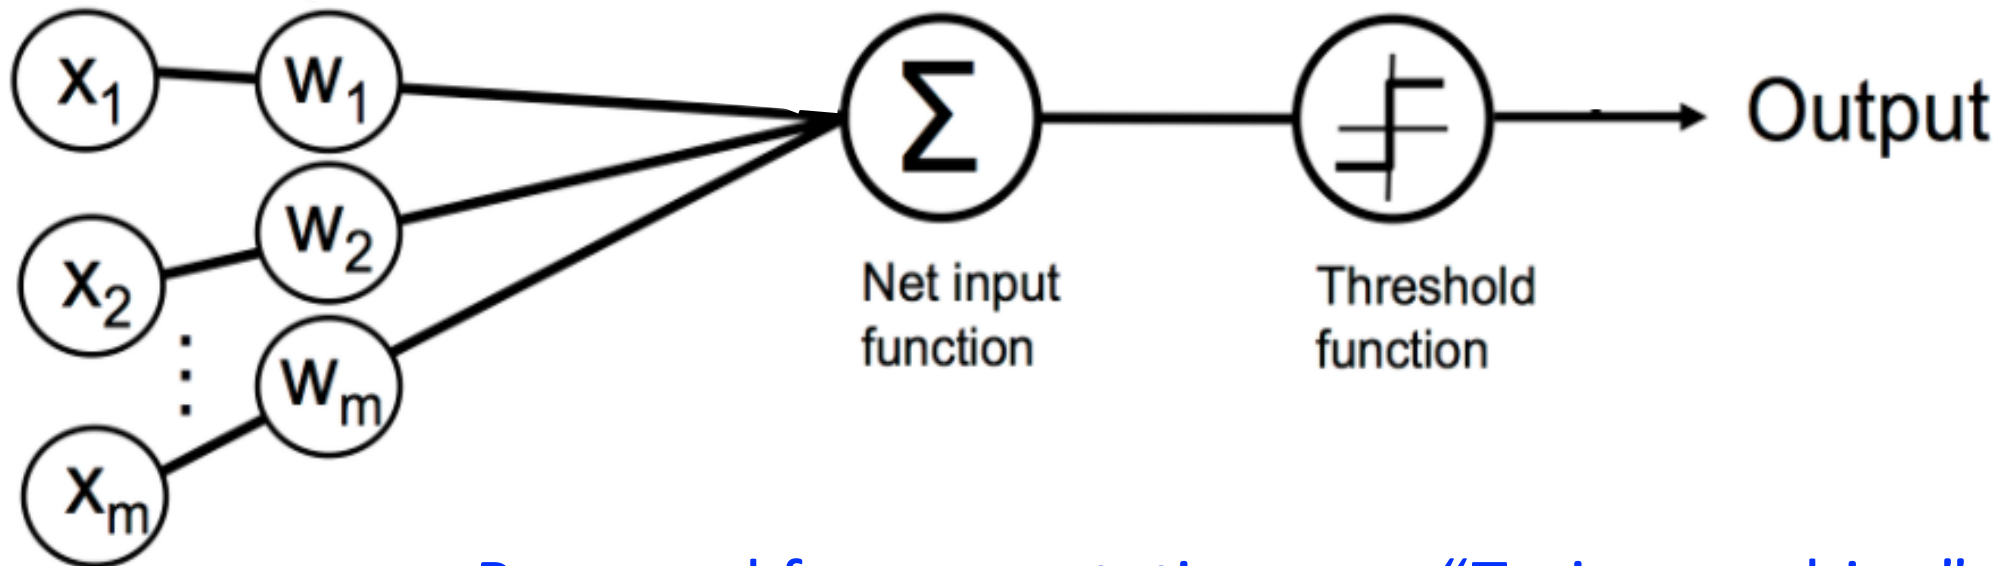
Artificial Neuron: McCulloch-Pitts Neuron

- inputs (x) and weights (w) can be 0 or 1
- weights (w) and threshold values are fixed
- outputs 1 or 0 (mimics neurons by “firing” only when aggregate value exceeds threshold)



Artificial Neuron: McCulloch-Pitts Neuron

- inputs (x) and weights (w) can be 0 or 1
- weights (w) and threshold values are fixed
- outputs 1 or 0 (mimics neurons by “firing” only when aggregate value exceeds threshold)



Proposed for computation on a “Turing machine”

Perceptron: Model Mimicking Human Machinery

First mathematical
model of neuron

1943 1945 1950




First programmable
machine

Turing test

1956
AI
Perceptron

1957
Machine
learning

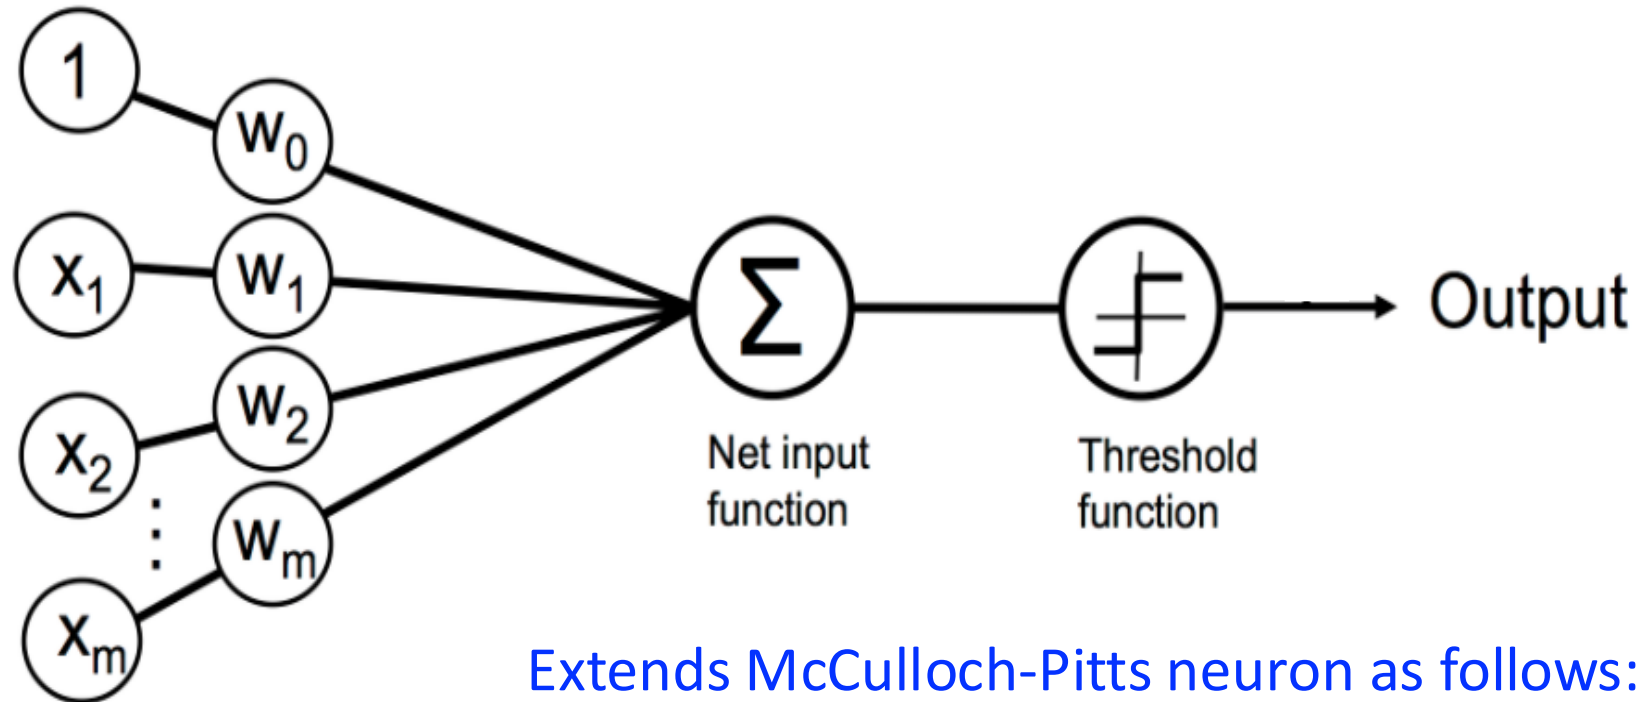


Frank Rosenblatt
(Psychologist)

<https://news.cornell.edu/stories/2019/09/professors-perceptron-paved-way-ai-60-years-too-soon>

Frank Rosenblatt, The perceptron, a perceiving and recognizing automaton Project Para. Cornell Aeronautical Laboratory, 1957

Perceptron: Architecture (Linear Threshold Unit)



Extends McCulloch-Pitts neuron as follows:

- inputs and weights can be any value
- weights (W) are learned

Perceptron: Architecture (Linear Threshold Unit)

- Function deciding output value (“fire” or not):

$$\phi(z) = \begin{cases} 1 & \text{if } z \geq \theta \\ -1 & \text{otherwise} \end{cases}$$


- Rewriting function:

$$\phi(z) = \begin{cases} 1 & \text{if } z \geq 0 \\ -1 & \text{otherwise} \end{cases}$$

*** Note: Kamath textbook offers two common conventions for Perceptrons of using two possible output values of $\{-1, 1\}$ and $\{0, 1\}$, in Chapters 2.5 and 4. The output choice dictates whether the threshold should be set to 0.5 or 0.**


- Where:

$$z = w_0 x_0 + w_1 x_1 + \dots + w_m x_m = \mathbf{w}^T \mathbf{x}$$

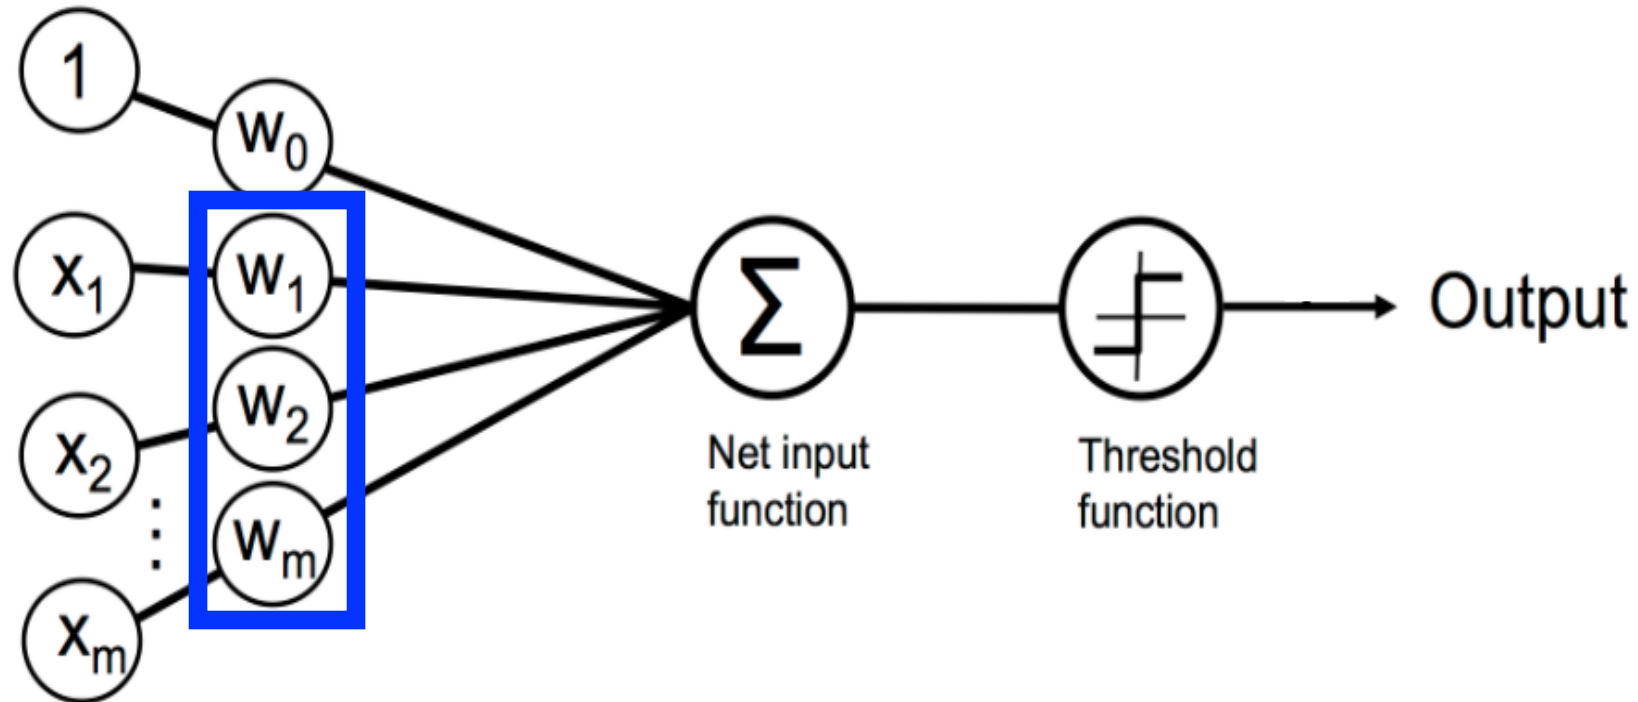


Perceptron: Architecture (Linear Threshold Unit)

Graphical representation
when there are two features:

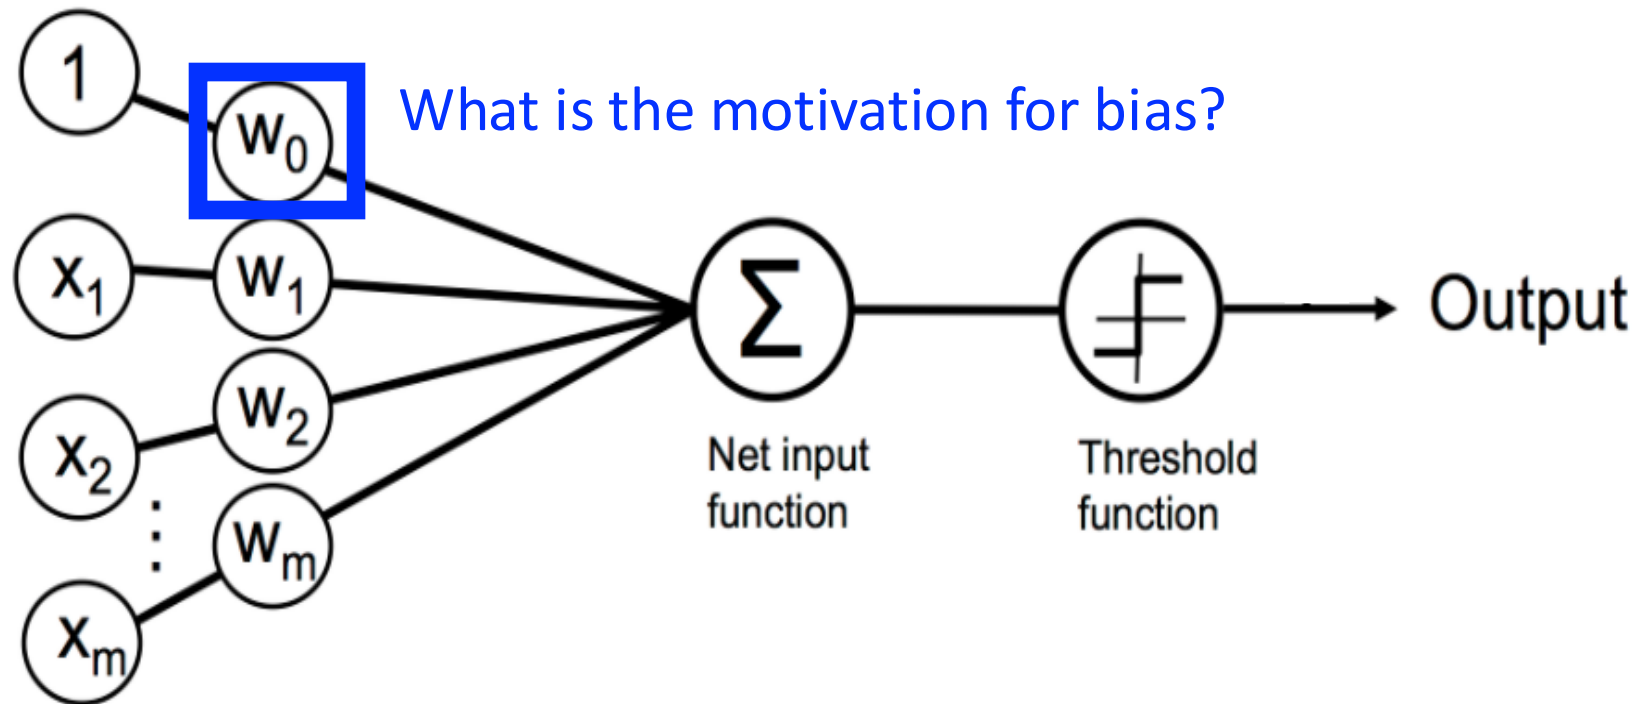


Perceptron: Architecture (Linear Threshold Unit)



What is the motivation for weights?
e.g., predicting if you will like a movie?

Perceptron: Architecture (Linear Threshold Unit)



Perceptron: Architecture (Linear Threshold Unit)

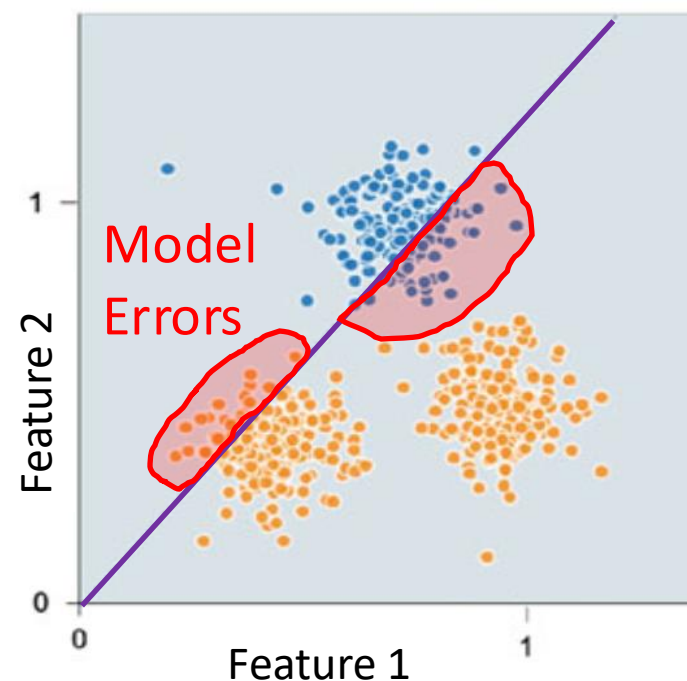
The model (in 2D, a line) must go through the origin without a **bias**:

$$z = w_0 x_0 + w_1 x_1 + \dots + w_m x_m = \mathbf{w}^T \mathbf{x}$$

Bias

$-\theta$	1
-----------	---

Binary classification problems
(separate blue and orange points):



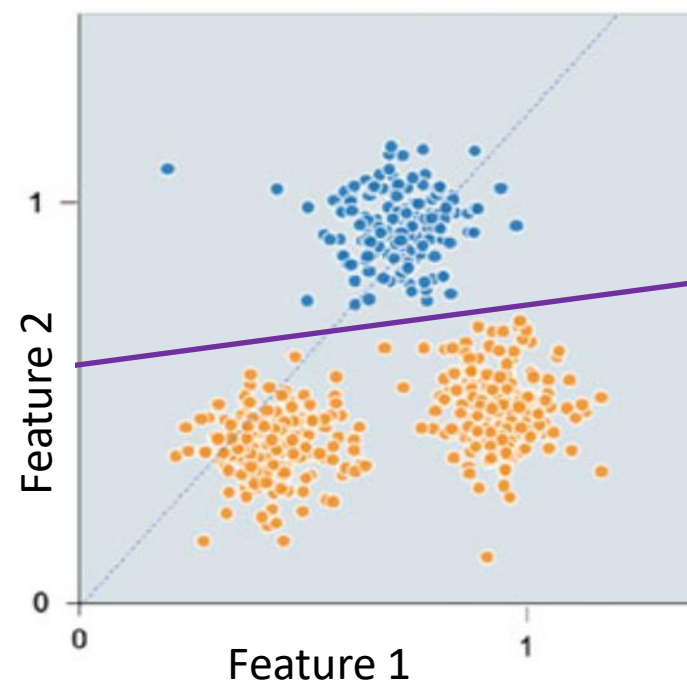
Perceptron: Architecture (Linear Threshold Unit)

The **model** (in 2D, a line) doesn't have to pass through the origin with **bias**:

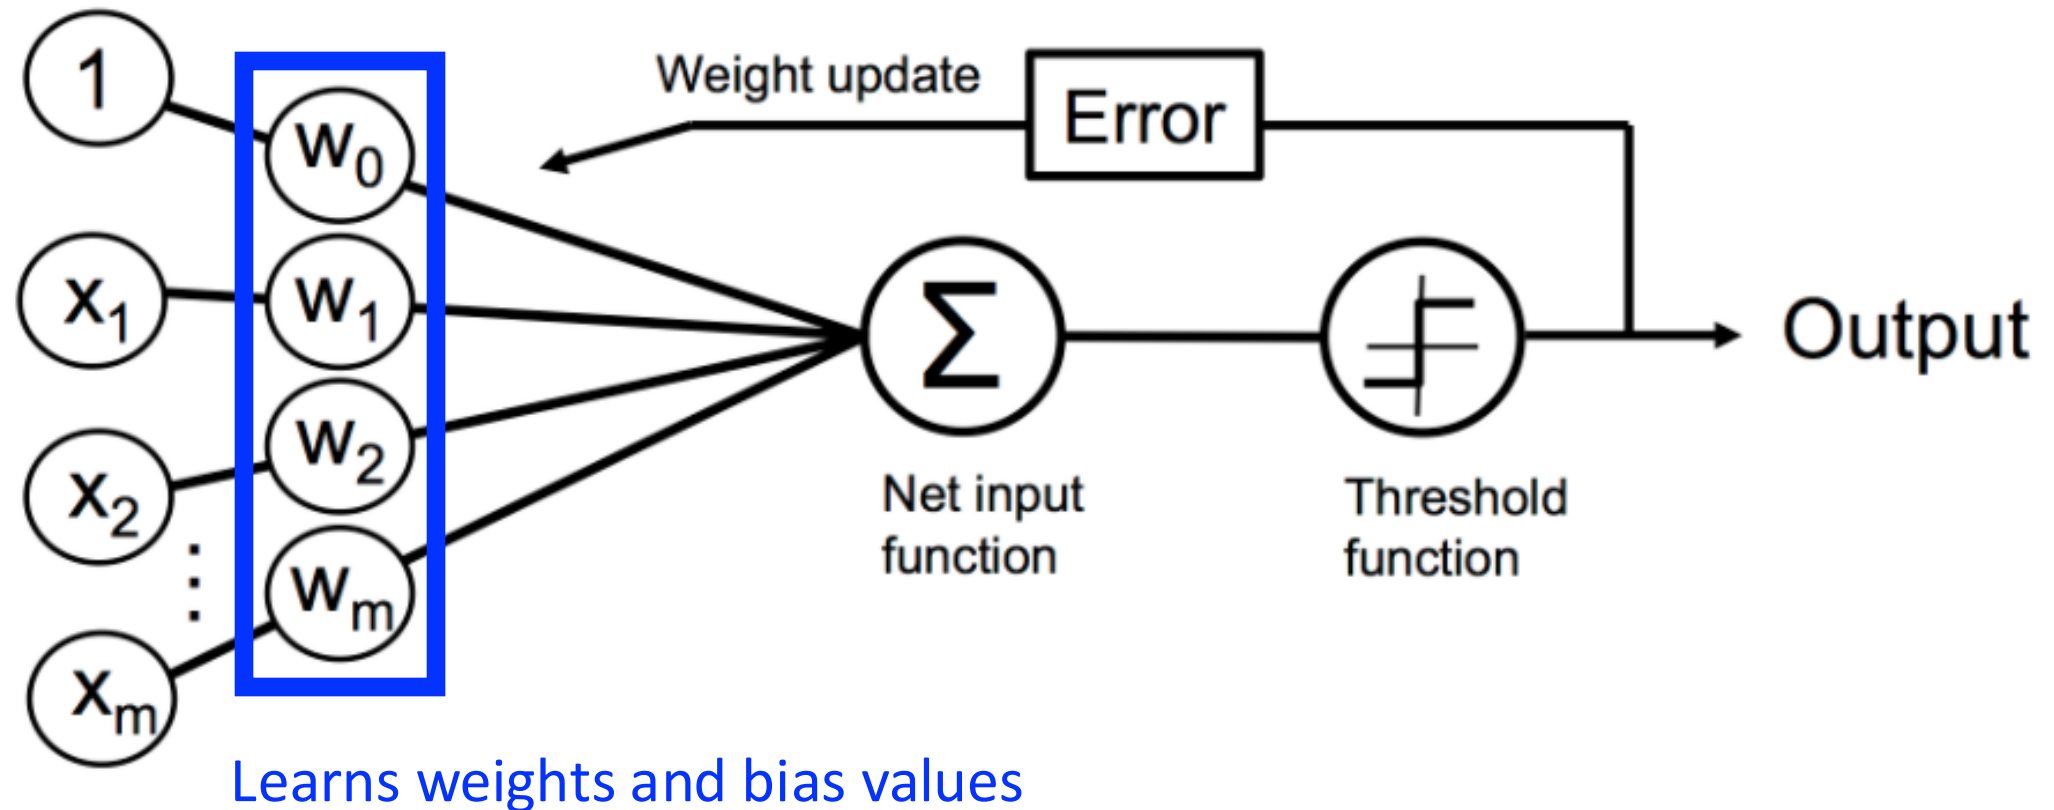
$$z = w_0 x_0 + w_1 x_1 + \dots + w_m x_m = \mathbf{w}^T \mathbf{x}$$

Bias $-\theta$ 1

Binary classification problems
(separate blue and orange points):




Perceptron: Learning Algorithm



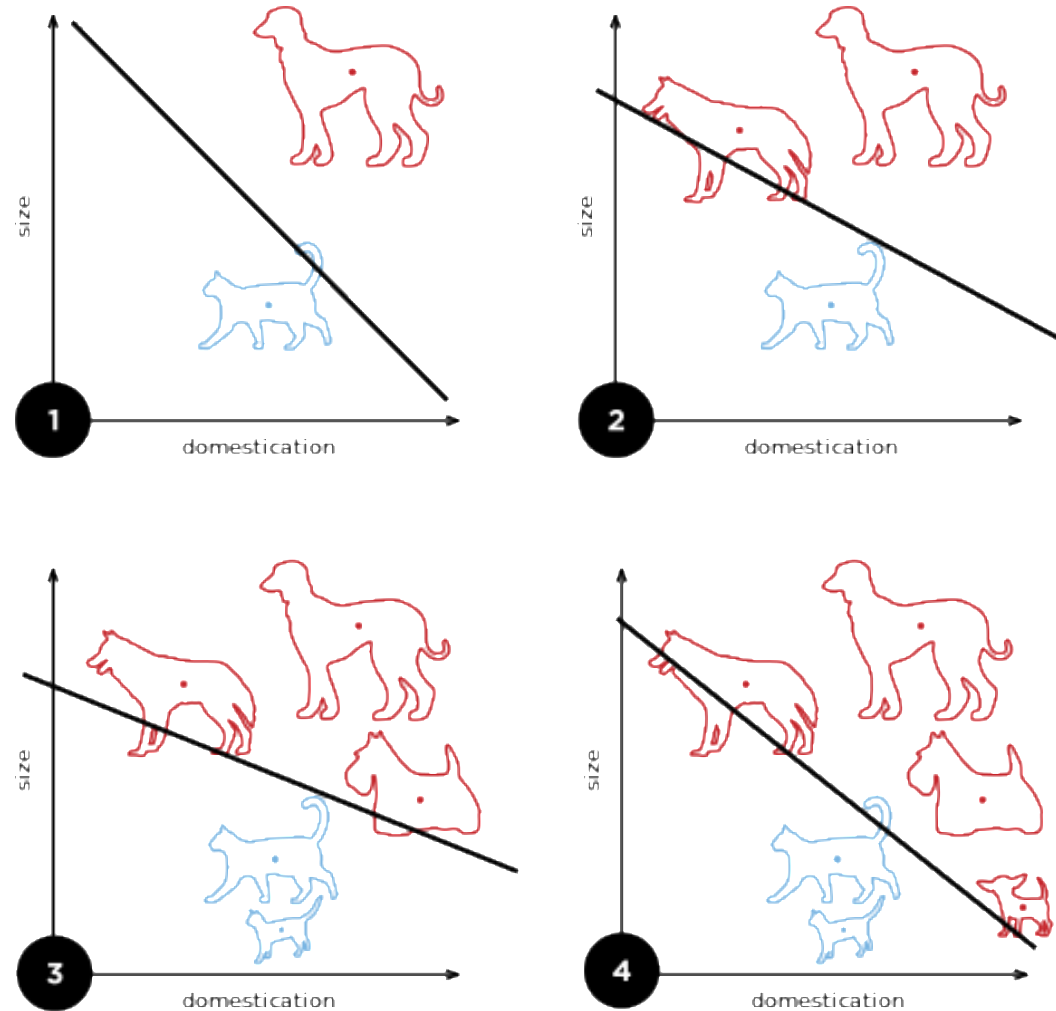
Perceptron: Learning Algorithm

Process: iteratively update boundary with observation of each additional example:



Perceptron: Learning Algorithm

Process: iteratively update boundary with observation of each additional example:



Perceptron: Learning Algorithm

1. Initialize weights/bias to 0 or small random numbers

2. Repeat until stopping criterion met:

1. Compute predicted value (i.e., $\{-1, 1\}$): $\sum_{j=0}^m \mathbf{x}_j \mathbf{w}_j = \mathbf{w}^T \mathbf{x}$

2. Update parameters based on prediction success: $w_j := w_j + \Delta w_j$.

$$\Delta w_j = \eta (\text{target}^{(i)} - \text{output}^{(i)}) x_j^{(i)}$$

(for i -th training example)

Learning Rate
(set a priori and held constant)

True Class Label

Predicted Class Label

Perceptron: Learning Algorithm

1. Initialize weights/bias to 0 or small random numbers

2. Repeat until stopping criterion met:

1. Compute predicted value (i.e., $\{-1, 1\}$): $\sum_{j=0}^m \mathbf{x}_j \mathbf{w}_j = \mathbf{w}^T \mathbf{x}$

2. Update parameters based on prediction success: $w_j := w_j + \Delta w_j$.

$$\Delta w_j = \eta (\text{target}^{(i)} - \text{output}^{(i)}) x_j^{(i)} \quad (\text{for } i\text{-th training example})$$

What happens to the parameters when the model predicts the **correct** class label?

- no update since result is 0

Perceptron: Learning Algorithm

1. Initialize weights/bias to 0 or small random numbers

2. Repeat until stopping criterion met:

1. Compute predicted value (i.e., $\{-1, 1\}$): $\sum_{j=0}^m \mathbf{x}_j \mathbf{w}_j = \mathbf{w}^T \mathbf{x}$

2. Update parameters based on prediction success: $w_j := w_j + \Delta w_j$.

$$\Delta w_j = \eta (\text{target}^{(i)} - \text{output}^{(i)}) x_j^{(i)} \quad (\text{for } i\text{-th training example})$$


What happens to the parameters when the model predicts the **wrong** class label?

- updates since result is "2" or "-2"

Perceptron: Example

- True Model: Y is 1 if at least 2 of the 3 inputs are 1, and -1 otherwise


X_1	X_2	X_3	Y
1	0	0	
1	0	1	
1	1	0	
1	1	1	?
0	0	1	
0	1	0	
0	1	1	
0	0	0	



Perceptron: Example

- True Model: Y is 1 if at least 2 of the 3 inputs are 1, and -1 otherwise


X_1	X_2	X_3	Y
1	0	0	-1
1	0	1	
1	1	0	
1	1	1	?
0	0	1	
0	1	0	
0	1	1	
0	0	0	



Perceptron: Example

- True Model: Y is 1 if at least 2 of the 3 inputs are 1, and -1 otherwise

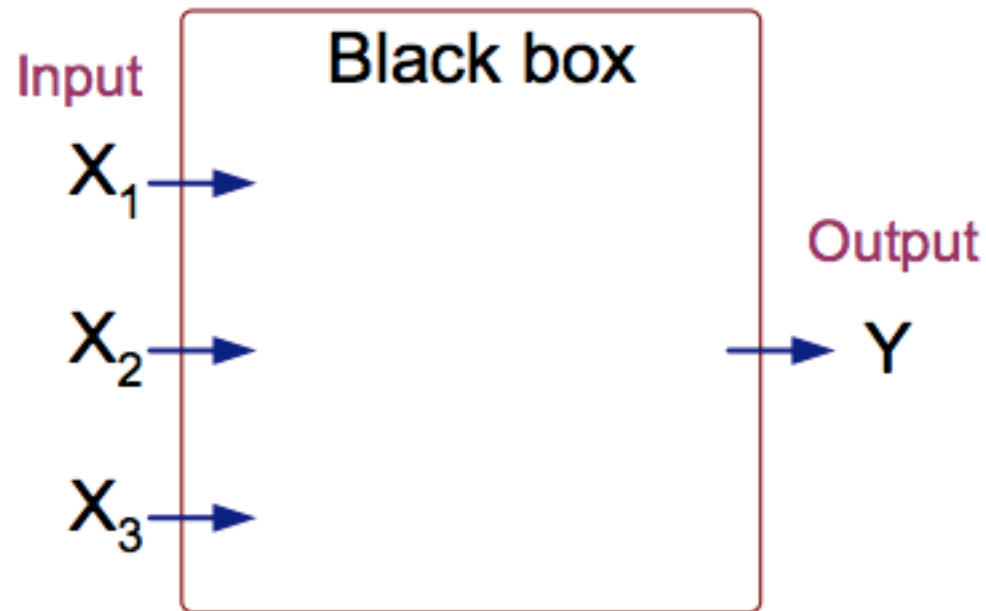
X_1	X_2	X_3	Y
1	0	0	-1
1	0	1	1
1	1	0	?
1	1	1	
0	0	1	
0	1	0	
0	1	1	
0	0	0	



Perceptron: Example

- True Model: Y is 1 if at least 2 of the 3 inputs are 1, and -1 otherwise

X_1	X_2	X_3	Y
1	0	0	-1
1	0	1	1
1	1	0	1
1	1	1	1
0	0	1	-1
0	1	0	-1
0	1	1	1
0	0	0	-1



Perceptron: Example (Training with 1st Sample)

- Compute predicted value: $\sum_{j=0}^m x_j w_j = \mathbf{w}^T \mathbf{x}$; $\phi(\mathbf{w}^T \mathbf{x}) = \begin{cases} 1 & \text{if } \phi(\mathbf{w}^T \mathbf{x}) \geq 0 \\ -1 & \text{otherwise} \end{cases}$

X_1	X_2	X_3	Y	Predicted	w_0	w_1	w_2	w_3
1	0	0	-1	?	0	0	0	0

Perceptron: Example (Training with 1st Sample)

- Update params: $w_j = w_j + \eta (\text{target}^{(i)} - \text{output}^{(i)}) x_j^{(i)}$; learning rate = 0.1

X_1	X_2	X_3	Y	Predicted
1	0	0	-1	1

w_0	w_1	w_2	w_3
0	0	0	0
?	?	?	?

$$\Delta w_0 = \eta (\text{target}^{(i)} - \text{output}^{(i)})$$

$$\Delta w_1 = \eta (\text{target}^{(i)} - \text{output}^{(i)}) x_1^{(i)}$$

$$\Delta w_2 = \eta (\text{target}^{(i)} - \text{output}^{(i)}) x_2^{(i)}$$

$$\Delta w_3 = \eta (\text{target}^{(i)} - \text{output}^{(i)}) x_3^{(i)}$$

Perceptron: Example (Training with 1st Sample)

- Update params: $w_j = w_j + \eta (\text{target}^{(i)} - \text{output}^{(i)}) x_j^{(i)}$; learning rate = 0.1

X_1	X_2	X_3	Y
1	0	0	-1

Predicted
1

w_0	w_1	w_2	w_3
0	0	0	0
?	?	?	?

$$\Delta w_0 = 0.1(-1-1)*1 = -0.2$$

$$\Delta w_1 = 0.1(-1-1)*1 = -0.2$$

$$\Delta w_2 = 0.1(-1-1)*0 = 0$$

$$\Delta w_3 = 0.1(-1-1)*0 = 0$$

When predicted value is greater than true value (e.g., $1 > -1$), the updates *decrease* parameter values to increase the likelihood of classifying the sample as -1 next time

Perceptron: Example (Training with 2nd Sample)

- Compute output value: $\sum_{j=0}^m x_j w_j = \mathbf{w}^T \mathbf{x}$; $\phi(\mathbf{w}^T \mathbf{x}) = \begin{cases} 1 & \text{if } \phi(\mathbf{w}^T \mathbf{x}) \geq 0 \\ -1 & \text{otherwise} \end{cases}$

X_1	X_2	X_3	Y
1	0	0	-1
1	0	1	1

Predicted
1
?

w_0	w_1	w_2	w_3
0	0	0	0
-0.2	-0.2	0	0

Perceptron: Example (Training with 2nd Sample)

- Update params: $w_j = w_j + \eta (\text{target}^{(i)} - \text{output}^{(i)}) x_j^{(i)}$; learning rate = 0.1

X_1	X_2	X_3	Y
1	0	0	-1
1	0	1	1

Predicted
1
-1

w_0	w_1	w_2	w_3
0	0	0	0
-0.2	-0.2	0	0
?	?	?	?

$$\Delta w_0 = \eta (\text{target}^{(i)} - \text{output}^{(i)})$$

$$\Delta w_1 = \eta (\text{target}^{(i)} - \text{output}^{(i)}) x_1^{(i)}$$

$$\Delta w_2 = \eta (\text{target}^{(i)} - \text{output}^{(i)}) x_2^{(i)}$$

$$\Delta w_3 = \eta (\text{target}^{(i)} - \text{output}^{(i)}) x_3^{(i)}$$

Perceptron: Example (Training with 2nd Sample)

- Update params: $w_j = w_j + \eta (\text{target}^{(i)} - \text{output}^{(i)}) x_j^{(i)}$; learning rate = 0.1

X_1	X_2	X_3	Y
1	0	0	-1
1	0	1	1

Predicted
1
-1

w_0	w_1	w_2	w_3
0	0	0	0
-0.2	-0.2	0	0
?	?	?	?

$$\Delta w_0 = 0.1(1 - -1) * 1 = 0.2$$

$$\Delta w_1 = 0.1(1 - -1) * 1 = 0.2$$

$$\Delta w_2 = 0.1(1 - -1) * 0 = 0$$

$$\Delta w_3 = 0.1(1 - -1) * 1 = 0.2$$

When predicted value is less than true value (e.g., $-1 < 1$), the updates *increase* parameter values to increase the likelihood of classifying the sample as 1 next time

Perceptron: Example (Training with 2nd Sample)

- Update params: $w_j = w_j + \eta (\text{target}^{(i)} - \text{output}^{(i)}) x_j^{(i)}$; learning rate = 0.1

X_1	X_2	X_3	Y	Predicted
1	0	0	-1	1
1	0	1	1	-1

w_0	w_1	w_2	w_3
0	0	0	0
-0.2	-0.2	0	0
0	0	0	0.2

$$\Delta w_0 = 0.1(1 - -1) * 1 = 0.2$$

$$\Delta w_1 = 0.1(1 - -1) * 1 = 0.2$$

$$\Delta w_2 = 0.1(1 - -1) * 0 = 0$$

$$\Delta w_3 = 0.1(1 - -1) * 1 = 0.2$$

What is the influence of the learning rate? i.e., what would happen if the value was larger/smaller?

Perceptron: Example – One Epoch (Training with All Samples)

- $w_j = w_j + \eta (\text{target}^{(i)} - \text{output}^{(i)}) x_j^{(i)}$; learning rate = 0.1

X_1	X_2	X_3	Y
1	0	0	-1
1	0	1	1
1	1	0	1
1	1	1	1
0	0	1	-1
0	1	0	-1
0	1	1	1
0	0	0	-1

	w_0	w_1	w_2	w_3
0	0	0	0	0
1	-0.2	-0.2	0	0
2	0	0	0	0.2
3	0	0	0	0.2
4	0	0	0	0.2
5	-0.2	0	0	0
6	-0.2	0	0	0
7	0	0	0.2	0.2
8	-0.2	0	0.2	0.2

Perceptron: Example – Six Epochs

- $w_j = w_j + \eta (\text{target}^{(i)} - \text{output}^{(i)}) x_j^{(i)}$; learning rate = 0.1

X_1	X_2	X_3	Y
1	0	0	-1
1	0	1	1
1	1	0	1
1	1	1	1
0	0	1	-1
0	1	0	-1
0	1	1	1
0	0	0	-1

Epoch	w_0	w_1	w_2	w_3
0	0	0	0	0

Perceptron: Example – Six Epochs

- $w_j = w_j + \eta (\text{target}^{(i)} - \text{output}^{(i)}) x_j^{(i)}$; learning rate = 0.1

X_1	X_2	X_3	Y
1	0	0	-1
1	0	1	1
1	1	0	1
1	1	1	1
0	0	1	-1
0	1	0	-1
0	1	1	1
0	0	0	-1

	w_0	w_1	w_2	w_3
0	0	0	0	0
1	-0.2	-0.2	0	0
2	0	0	0	0.2
3	0	0	0	0.2
4	0	0	0	0.2
5	-0.2	0	0	0
6	-0.2	0	0	0
7	0	0	0.2	0.2
8	-0.2	0	0.2	0.2

Epoch	w_0	w_1	w_2	w_3
0	0	0	0	0
1	-0.2	0	0.2	0.2
2	-0.2	0	0.4	0.2
3	-0.4	0	0.4	0.2
4	-0.4	0.2	0.4	0.4
5	-0.6	0.2	0.4	0.2
6	-0.6	0.4	0.4	0.2

Perceptron: Learning Algorithm Choices

1. Initialize weights and bias **1. Values**
2. Repeat until stopping criterion met: **2.**

1. Compute predicted value (i.e., $\{-1, 1\}$): $\sum_{j=0}^m \mathbf{x}_j \mathbf{w}_j = \mathbf{w}^T \mathbf{x}$

2. Update parameters based on prediction success: $w_j := w_j + \Delta w_j$.

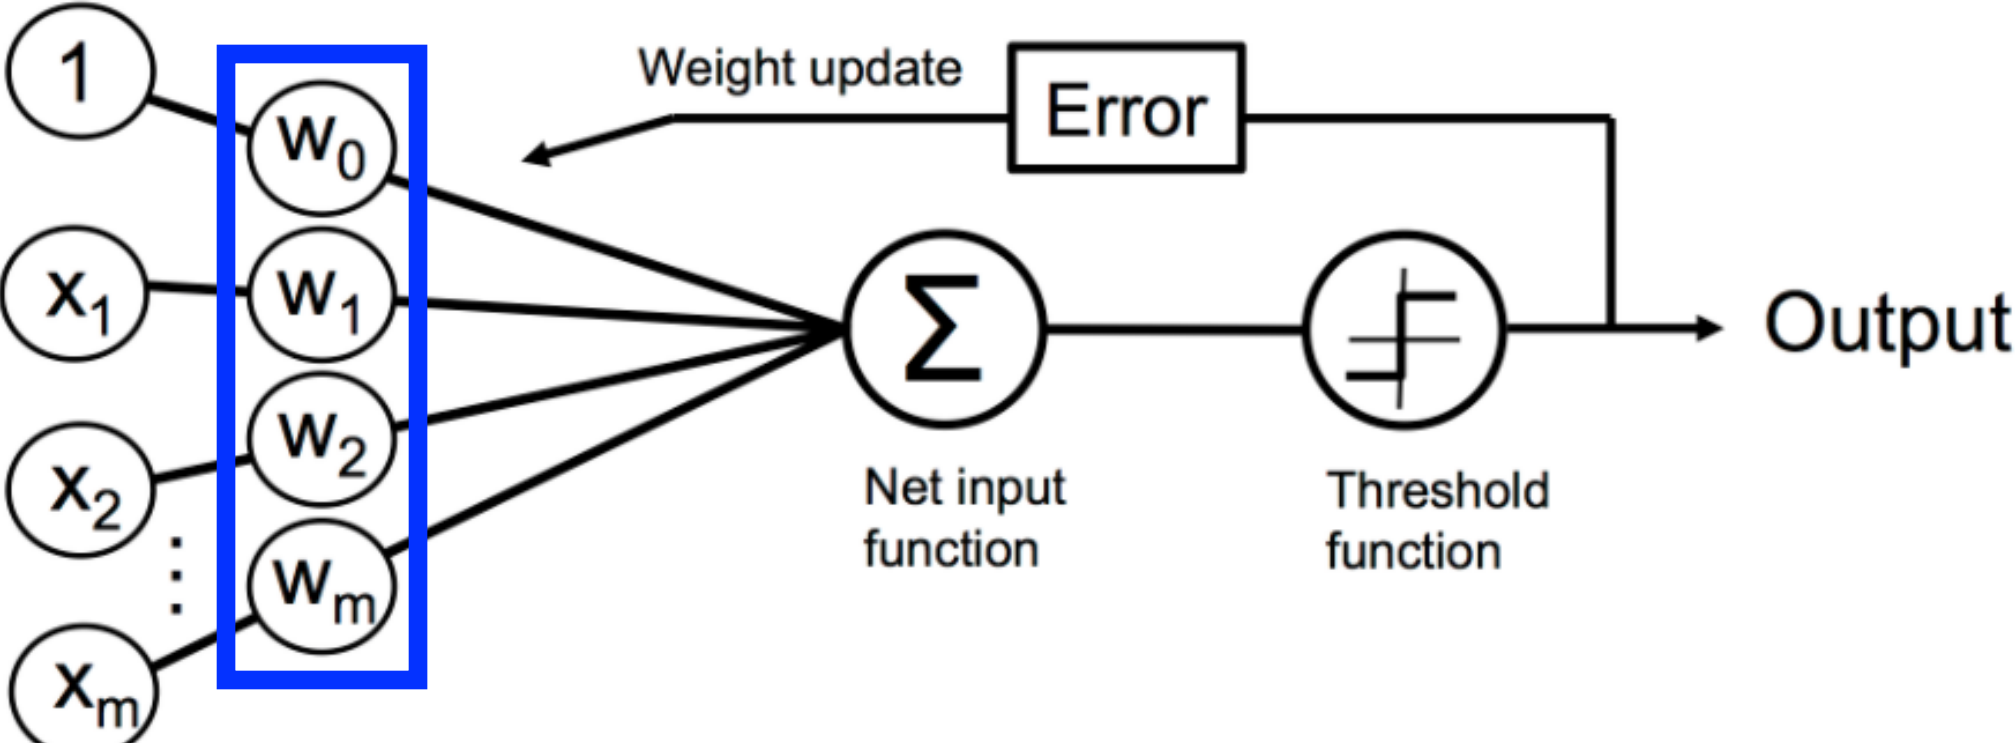
$$\Delta w_j = \eta (\text{target}^{(i)} - \text{output}^{(i)}) x_j^{(i)}$$

3. Learning Rate

Today's Topics

- Supervised learning: approach to develop a model
- Artificial neuron model: basic unit of neural networks
- Evaluating classification models

Evaluation: How Good Is Our Model?



Learned values for weights and bias

Evaluation Methods: Confusion Matrix

		Actual	
		Hairy	Not hairy
Predicted	Hairy	TP	FP
	Not hairy	FN	TN

TP = true positive

TN = true negative

FP = false positive

FN = false negative

Evaluation Methods: Descriptive Statistics

Commonly-used statistical descriptions:

e.g.,

		Actual	
		Hairy	Not hairy
Predicted	Hairy	50	10
	Not hairy	15	100

- How many **actual hairy** results are there? - 65
- How many **actual not hairy** results are there? - 110
- How many **correctly classified instances**? - 150/175 ~ 86%
- How many **incorrectly classified instances**? - 25/175 ~ 14%
- What is the **precision (aka – sensitivity)**?

(proportion of predicted instances actually hairy)

- $50/(50+10) \sim 83\%$

$$\frac{TP}{TP + FP}$$

- What is the **recall**? (proportion of hairy instances predicted as hairy)

- $50/(50+15) \sim 77\%$

$$\frac{TP}{TP + FN}$$

Evaluation Methods: Descriptive Statistics

(a)

		Actual	
		Z	Not Z
Predicted	Z	100	20
	Not Z	20	100

(b)

		Actual	
		Z	Not Z
Predicted	Z	80	0
	Not Z	120	40

(c)

		Actual	
		Z	Not Z
Predicted	Z	100	100
	Not Z	0	40

Which confusion matrix reflects the best-performing model?

Evaluation Methods: Descriptive Statistics

(a)

		Actual	
		Z	Not Z
Predicted	Z	100	20
	Not Z	20	100

(b)

		Actual	
		Z	Not Z
Predicted	Z	80	0
	Not Z	120	40

A *cautious* model that is correct whenever it predicts a positive label

(c)

		Actual	
		Z	Not Z
Predicted	Z	100	100
	Not Z	0	40

However, it can incorrectly predict many positives have negative labels

Which confusion matrix reflects the model with the highest precision?

$$\frac{TP}{TP + FP}$$

Evaluation Methods: Descriptive Statistics

(a)

		Actual	
		Z	Not Z
Predicted	Z	100	20
	Not Z	20	100

(b)

		Actual	
		Z	Not Z
Predicted	Z	80	0
	Not Z	120	40

(c)

		Actual	
		Z	Not Z
Predicted	Z	100	100
	Not Z	0	40

However, many negatives can be deemed positive

An *ambitious* model that tries to identify every positive label

Which confusion matrix reflects the model with the highest recall?

$$\frac{TP}{TP + FN}$$

Class Discussion

- Which of these evaluation metrics would you use versus not use and why?
 - Accuracy (percentage of correctly classified examples)
 - Precision (percentage of relevant instances among retrieved instances)
 - Recall (percentage of relevant instances retrieved)
- Scenario 1: Medical test for a rare disease affecting one in every million people.
- Scenario 2: Deciding which emails to flag as spam.

Today's Topics

- Supervised learning: approach to develop a model
- Artificial neuron model: basic unit of neural networks
- Evaluating classification models

The image features a central area with a radial gradient background, transitioning from a light center to a darker outer edge. This central area is framed by a dark grey border that mimics the appearance of a film strip, with white rectangular sprocket holes along the top and bottom edges. The text "The End" is centered within this frame.

The End