CS 232: Artificial Intelligence

Fall 2026

Prof. Carolyn Anderson
Wellesley College

The New York Times

Global Rush to Fix Security Flaw Uncovered by A.I. Leaves Sites Offline

Governments and businesses in Kenya, the Netherlands and elsewhere were scrambling to prevent hackers from exploiting a vulnerability in PHP, a widely used programming language.

Listen · 3:05 min

Share full article

0



Reminders

- HW 6 due Nov. 2nd
- My help hours: today 3-4 in W423
- New policy: attend an AI-related research talk and submit a summary for your choice of:
  - a bonus late day
  - a class participation credit

VOTER-CENTERED ELECTORAL REFORM

Thursday, October 29th
12:45-2:10pm
SCI H105

Lunch is provided!

Short talks and a panel Q&A on electoral reforms in the U.S. and how academic work supports them. Come learn about exciting reforms like proportional representation, participatory budgeting, and ranked choice voting. Speakers are:



Rachel Hutchinson

Deputy Director of Research & Policy at FairVote



Renato Castelo

Director of the Office of Participatory Budgeting, City of Boston



Brian Brubach

Assistant Professor of Computer Science at Wellesley College

Organized By

Prof. Brian Brubach, Computer Science

Sponsored By

Hillary Rodham Clinton Center for Citizenship, Leadership, and Democracy

Institute for Mathematics and Democracy

AIT - BUDAPEST
COMPUTER & DATA SCIENCE

WELLESLEY
COMPUTER SCIENCE

TECH TALK

Emerging Technologies in Mobile Software Development

The Shift
How AI impacts coding, testing, documentation, and deployment.

The Reality
Balancing rapid prototyping against technical debt and security.

The Core Question
AI writes code faster than ever—who's checking it?



Hosted by
Péter Ekler who is an Associate Professor at BME, Co-Academic Director of AIT-Budapest, and an AI strategist at Zenitech. He brings over 15 years of academic and industry leadership in Applied AI and software development.

Come and join the talk. Lunch will be served!

LOCATION
SCI-H105

DATE
October 26

TIME
1 - 2 PM

www.ait-budapest.com

?:sb129 accessibility@wellesley.edu exp: 10/27

HELP US STUDY AI SUPPORTED TEXT ANALYSIS

60 mins

LOCATION: VIRTUAL



SIGN UP HERE!

WHAT YOU WILL DO?

Read short **K-drama** discussion excerpts, **identify ideas** and **themes**, try three kinds of **AI support**, and answer questions about your experience. Some tasks will be completed without AI. Participants will receive a $20 gift card upon completion of the study.

WHO CAN TAKE PART?

- Current Wellesley College student, age 18 or older.
- Comfortable reading and responding in English.
- Must not have independently completed a project involving coding qualitative data (introductory coursework or brief practice is allowed.)
- No K-drama knowledge required.
- Must have access to a laptop and an Internet connection.

Note: Taking part is voluntary. You may stop at any time. Your decision will not affect your grades, academic standing, employment, or access to College resources.

WELLESLEY

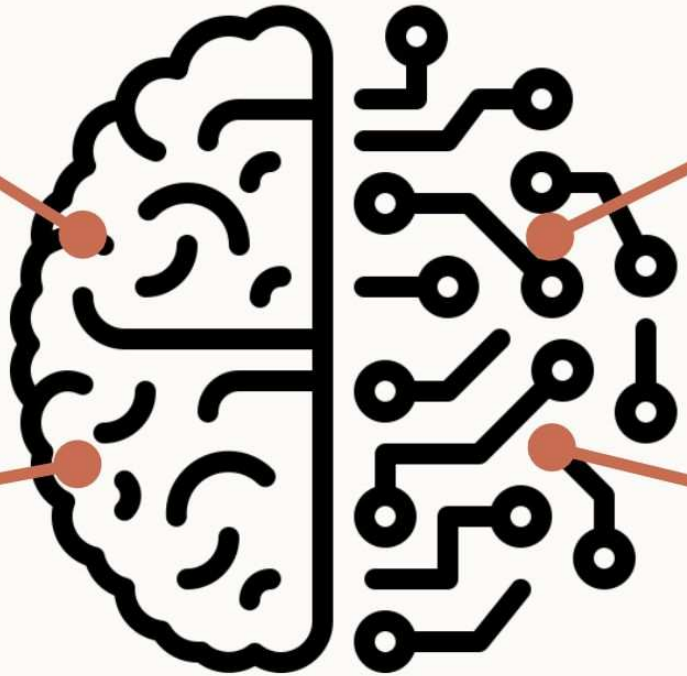
Questions @ hn103 | accessibility@wellesley.edu | exp 11/20/2026

Learn about the Cognitive and Linguistic Sciences Major

*The Science of **The Mind**, **Language**, **Computation**, and **Meaning***

Linguistics

Philosophy



Computer Science

Psychology

RSVP

Learn about our Spring '27 courses

Tuesday, Nov. 3rd 5-6pm SCI-H305

Snacks from Tatte
Meet Faculty & Majors

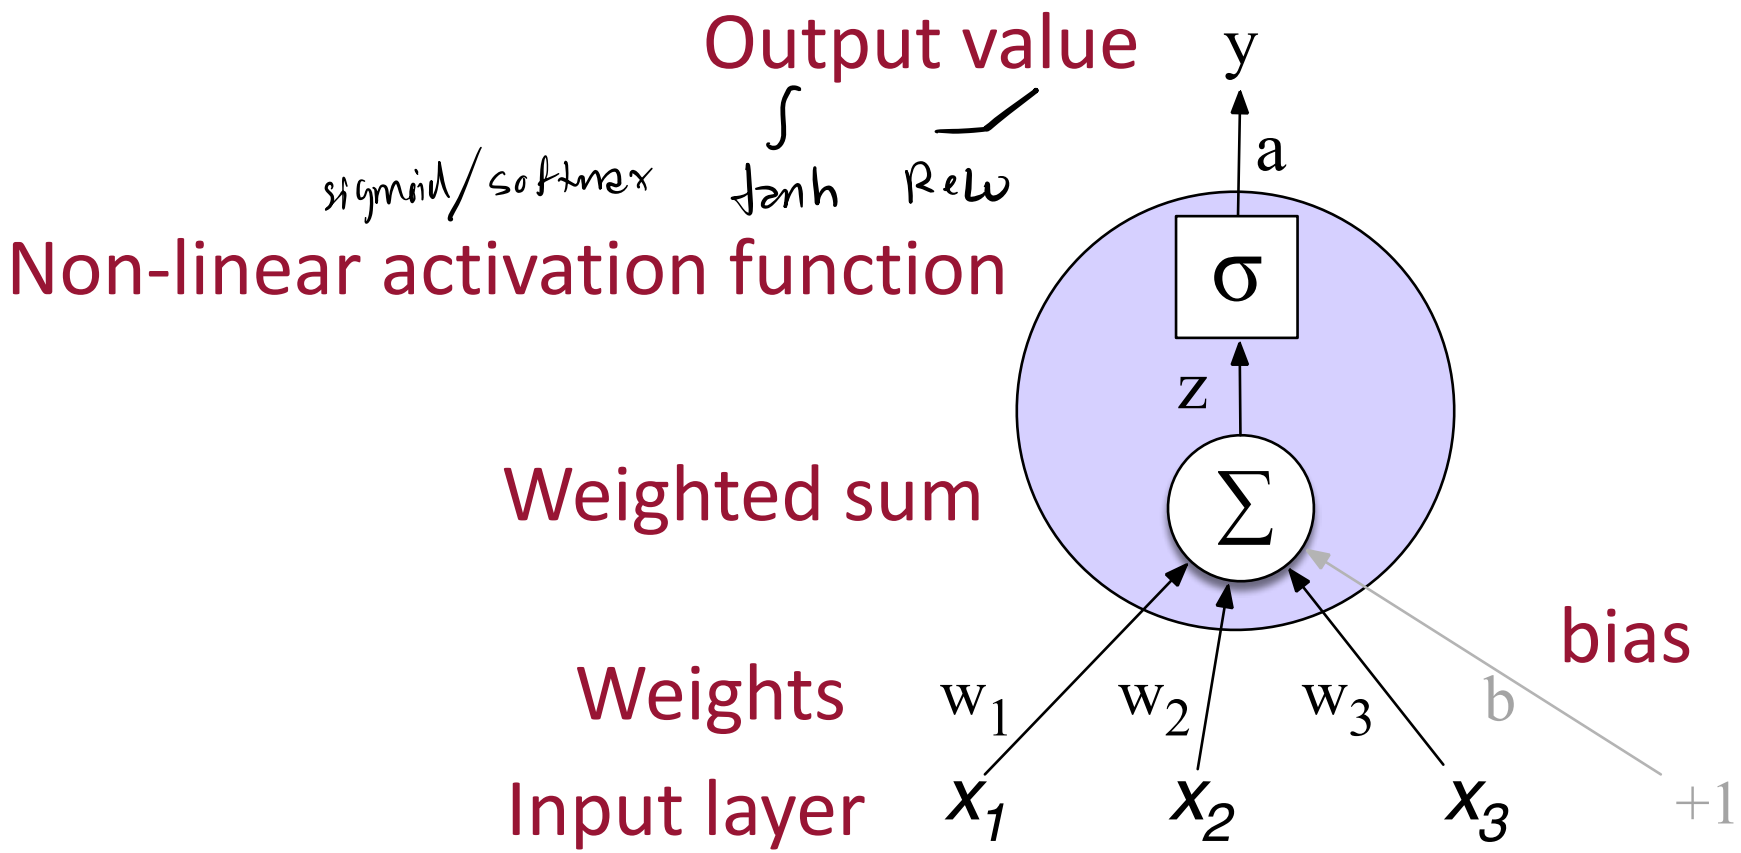
AI Needs You Discussion

Technological inevitability is the belief that it is futile to oppose technological change because it will happen anyway. Is this view consistent with the history of embryo research discussed in this chapter?

Recap: Neural Networks

Neural Network Unit

This is not in your brain



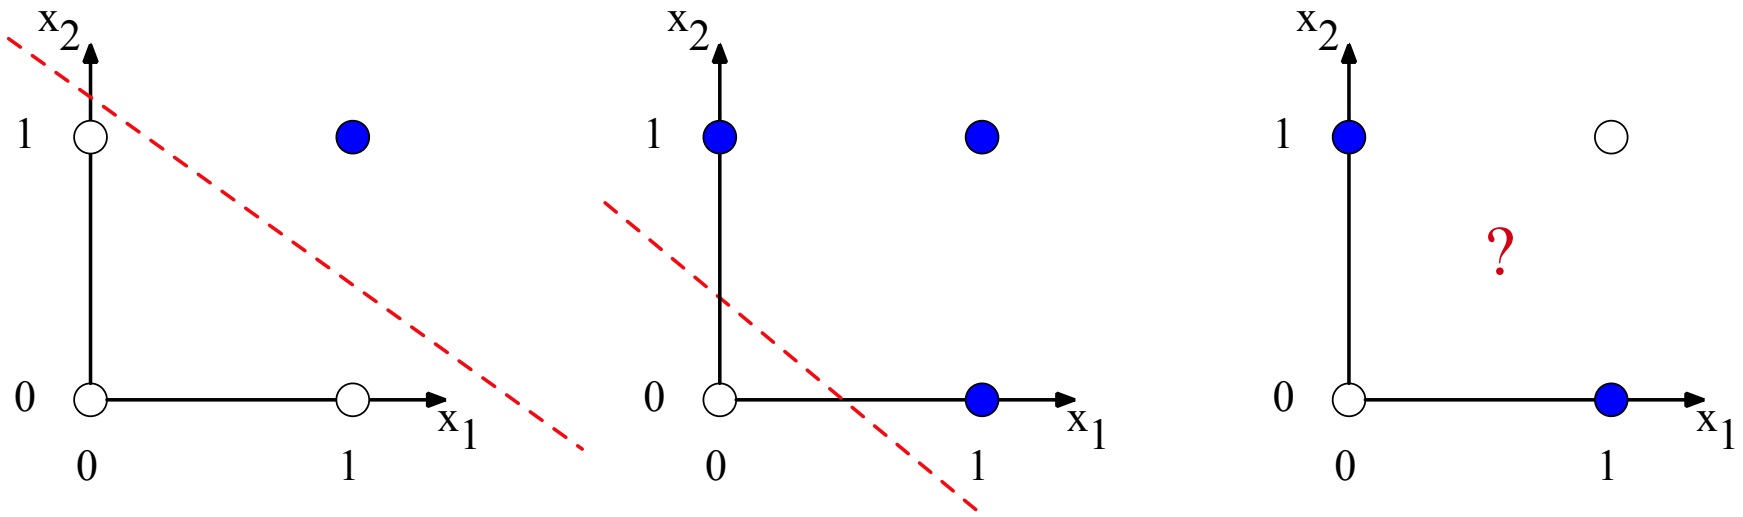
The XOR problem

Minsky and Papert (1969)

Can neural units compute simple functions of input?

✓ AND			✓OR			XOR		
x1	x2	y	x1	x2	y	x1	x2	y
0	0	0	0	0	0	0	0	0
0	1	0	0	1	1	0	1	1
1	0	0	1	0	1	1	0	1
1	1	1	1	1	1	1	1	0

Decision boundaries



a) x_1 AND x_2 b) x_1 OR x_2 c) x_1 XOR x_2

XOR is not a **linearly separable** function!

Solution to the XOR problem

XOR can't be calculated by a single perceptron

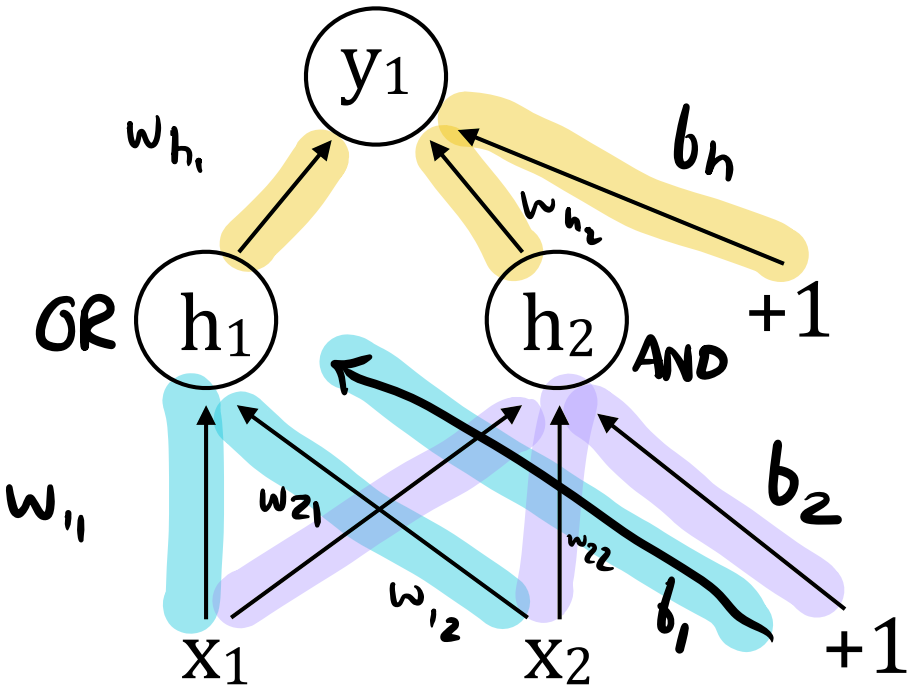
XOR can be calculated by a layered network of units.

XOR		
x1	x2	y
0	0	0
0	1	1
1	0	1
1	1	0

$w_{11} = 1$
$w_{12} = 1$
$b_1 = 0$

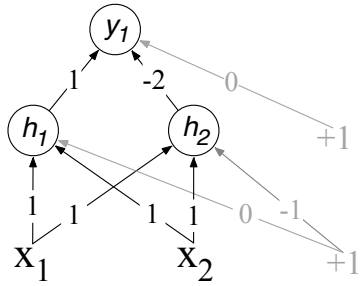
$w_{21} = 1$
$w_{22} = 1$
$b_2 = -1$

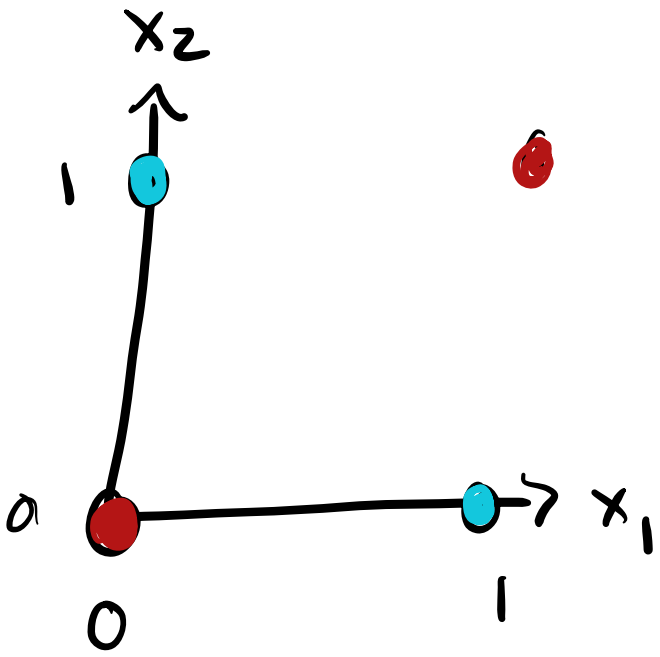
$w_{h_1} = 1$
$w_{h_2} = -2$
$b_h = 0$



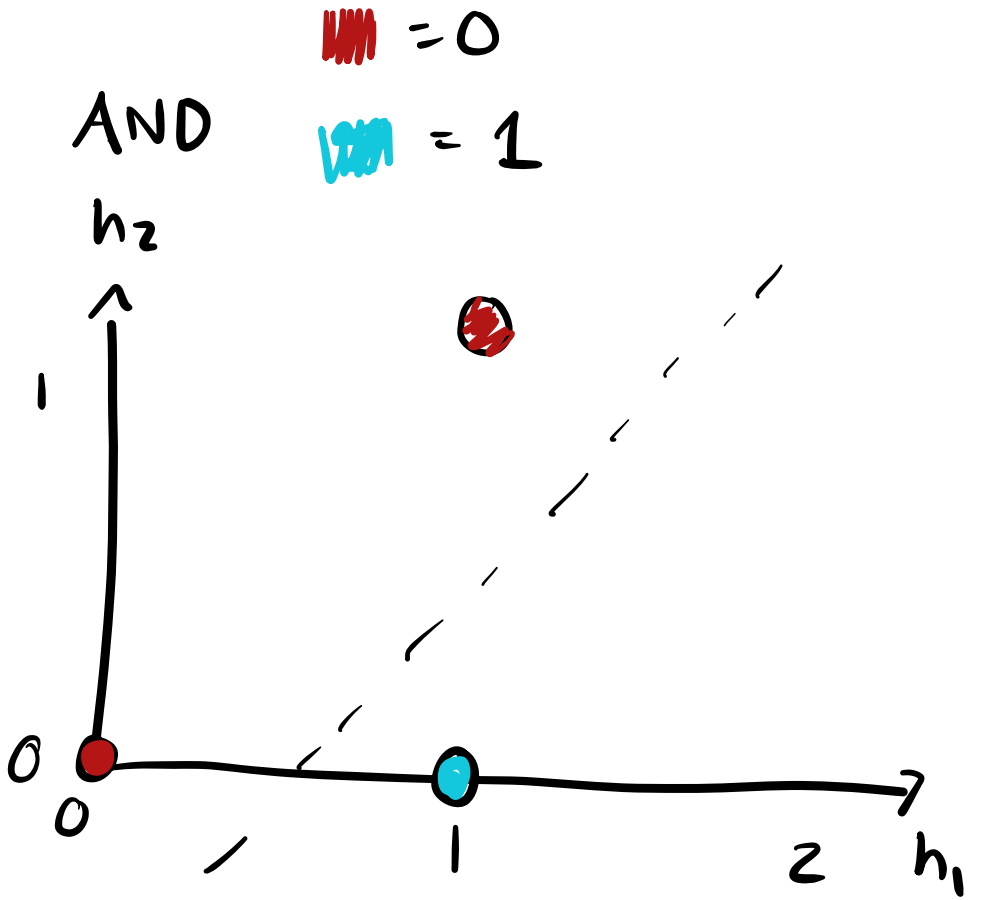
$$h_2 = \text{ReLu}(w_{21} x_1 + w_{22} x_2 + 1)$$

The hidden representation h





The original feature space



The new linearly seperable h space

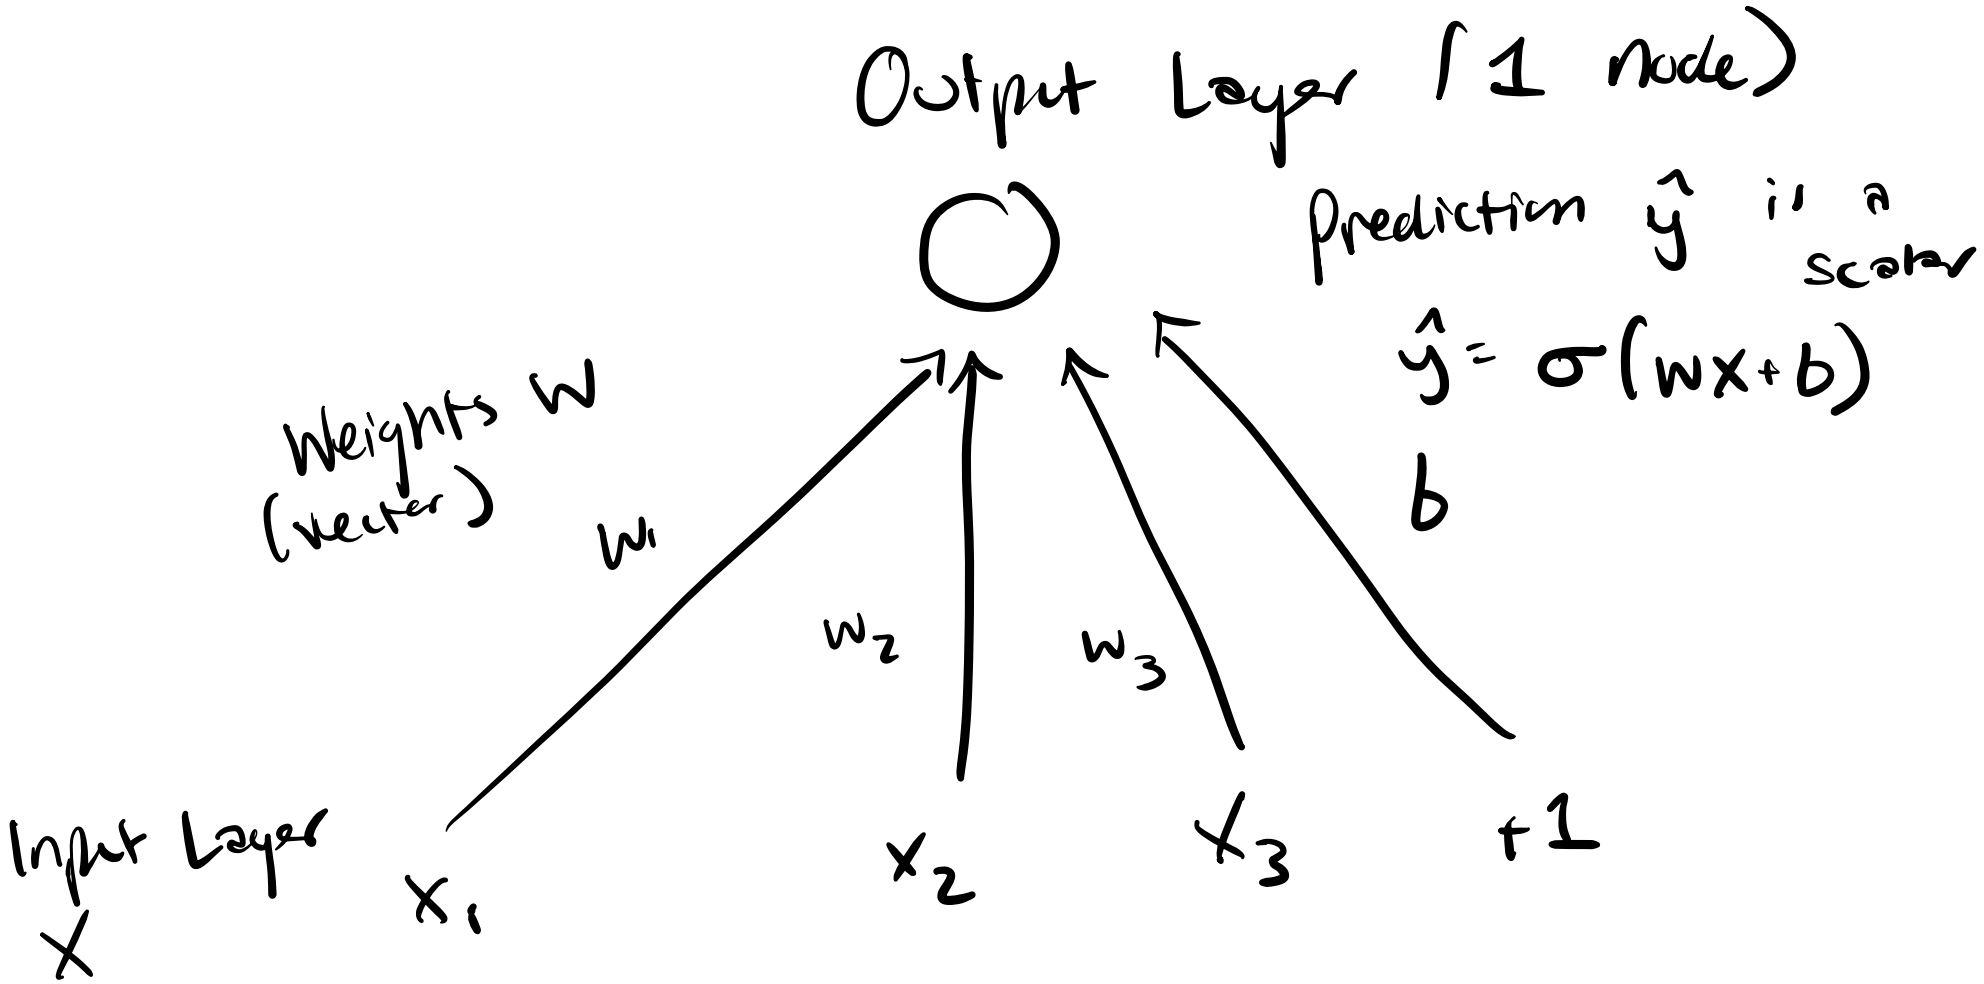
OR

h = hidden

Feedforward Networks

Binary Logistic Regression as a 1-layer Network

(we don't count the input layer when counting layers!)



Multinomial Logistic Regression as a 1-layer Network

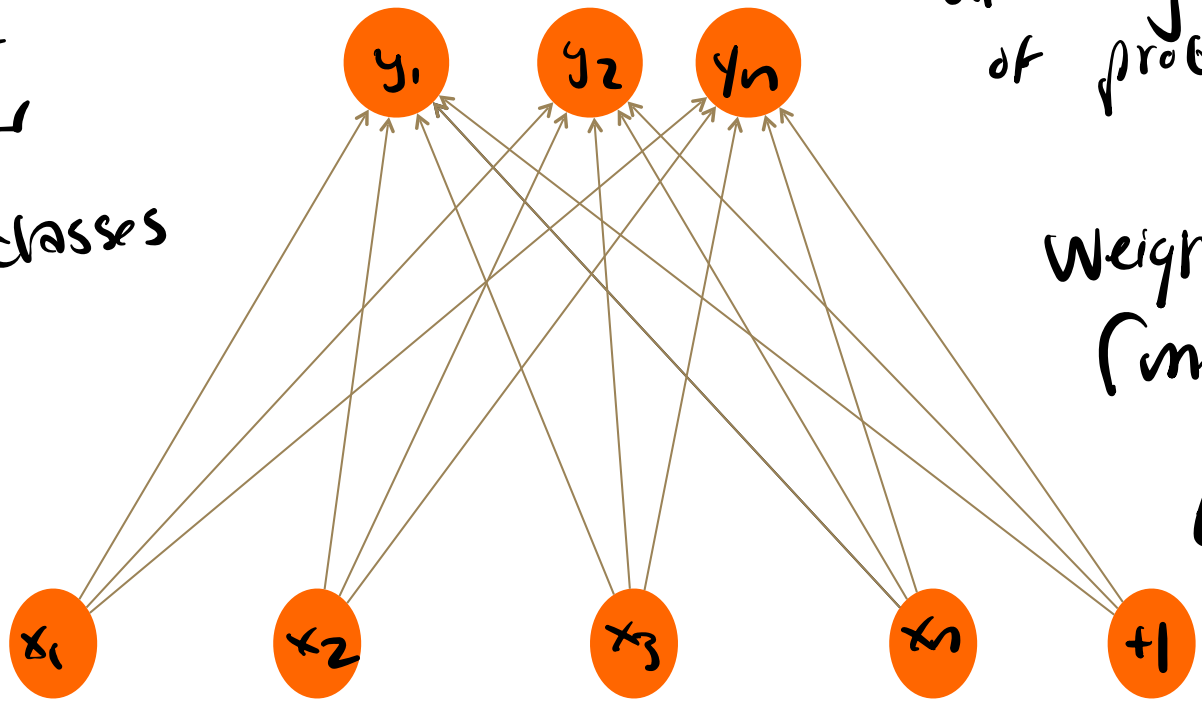
Fully connected single layer network

$\hat{y} = \text{softmax}(Wx + b)$

where $\hat{y}$ is a vector of probabilities

Output Layer

n nodes = n classes

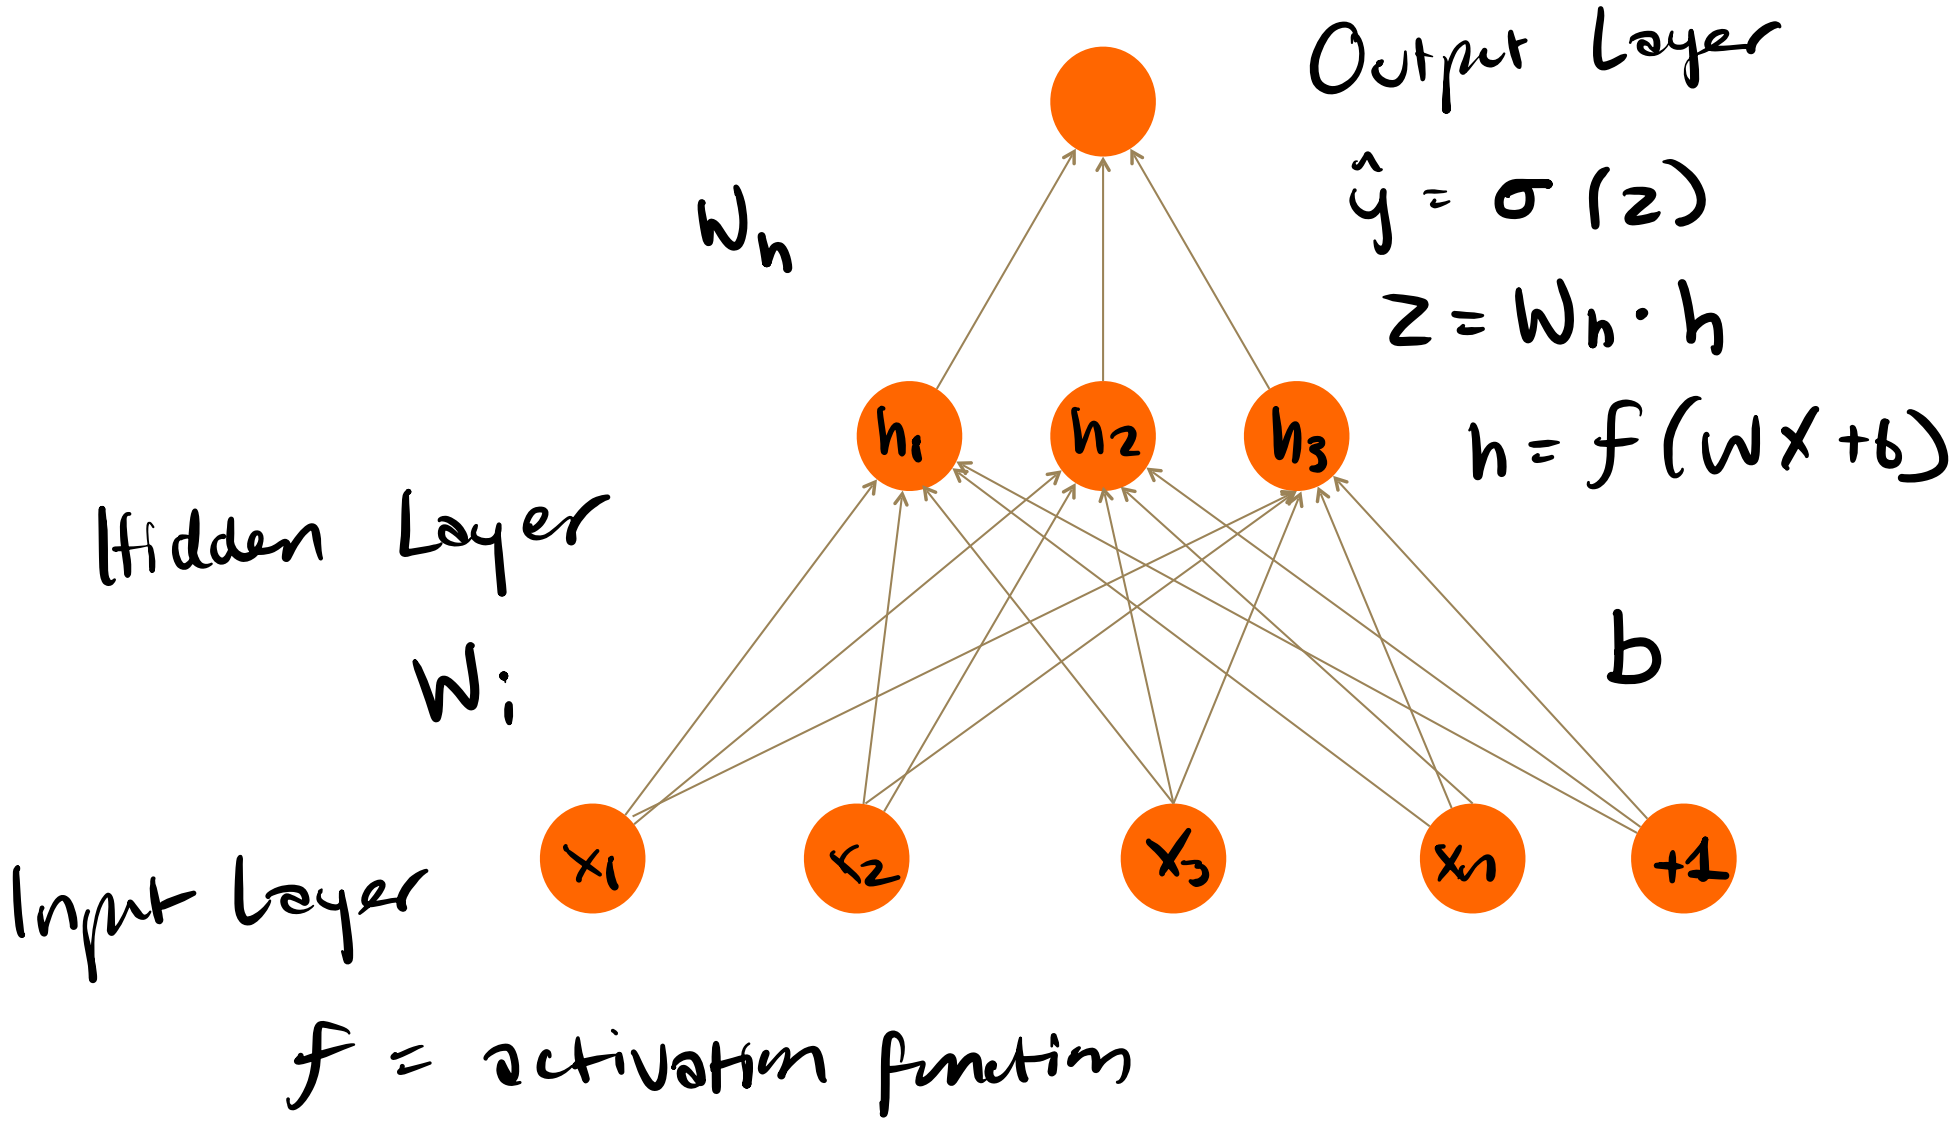


Weights W (matrix)

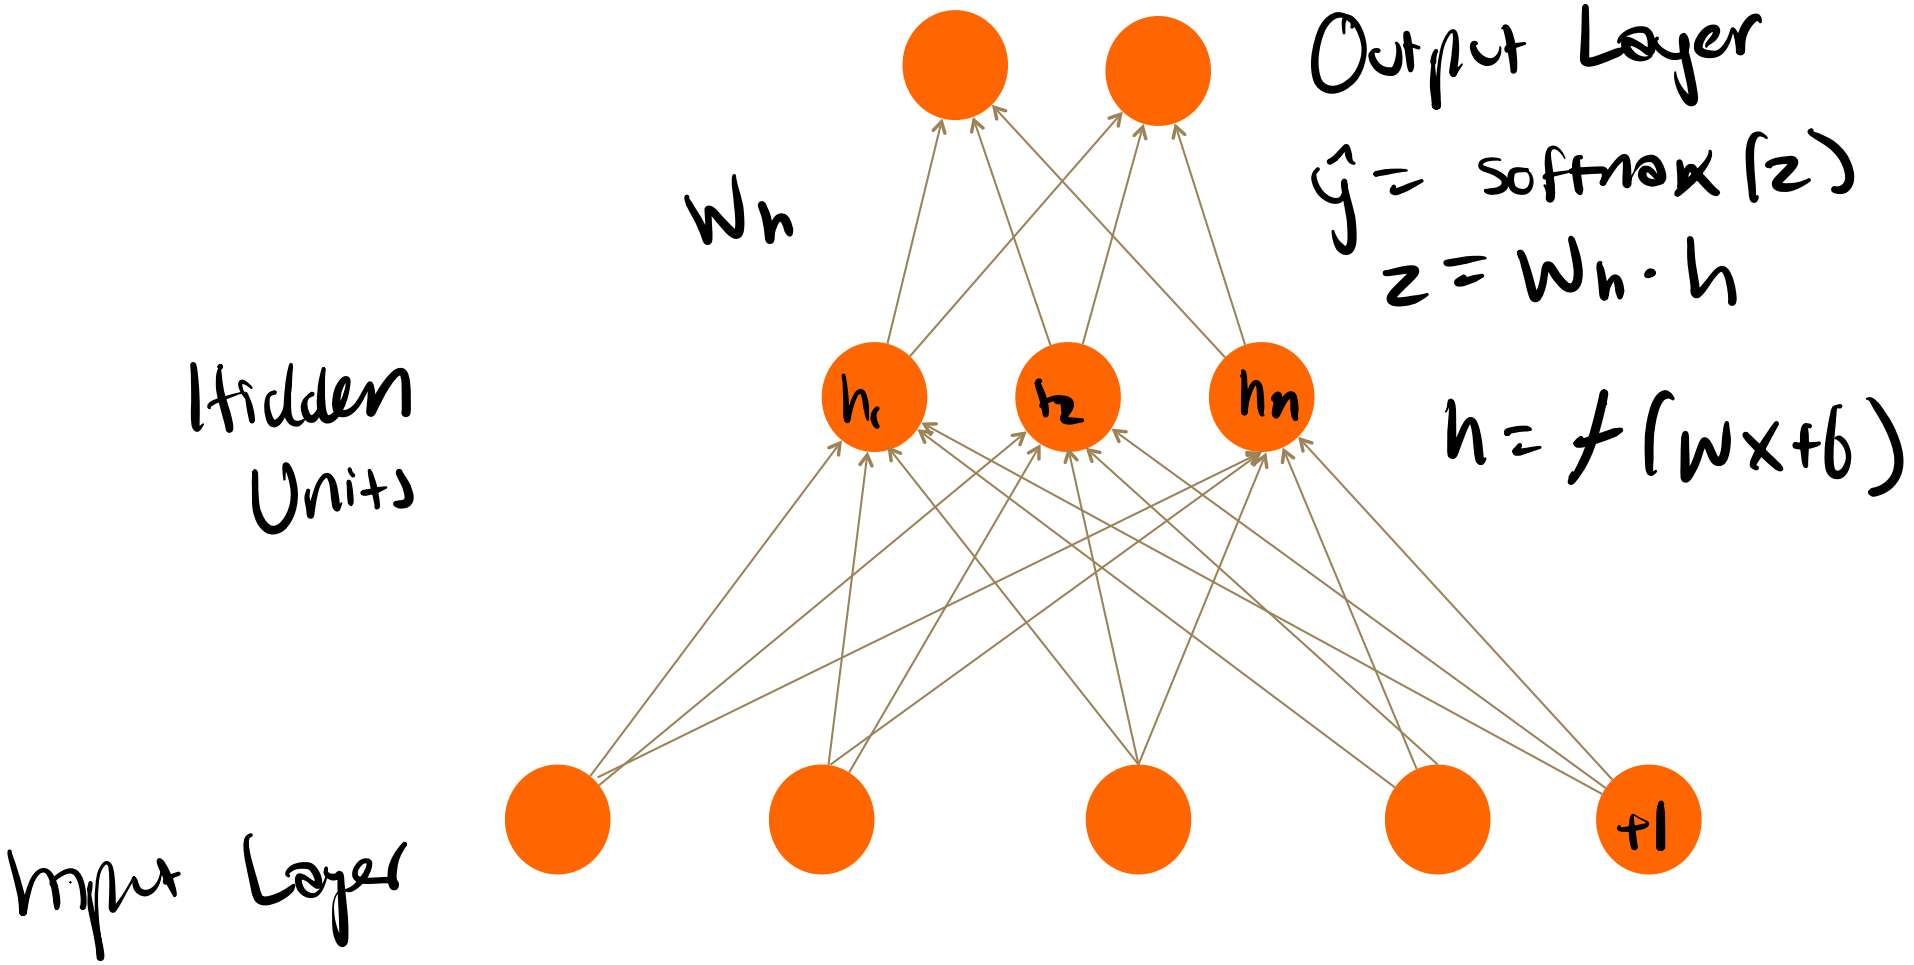
b (vector)

Input Layer

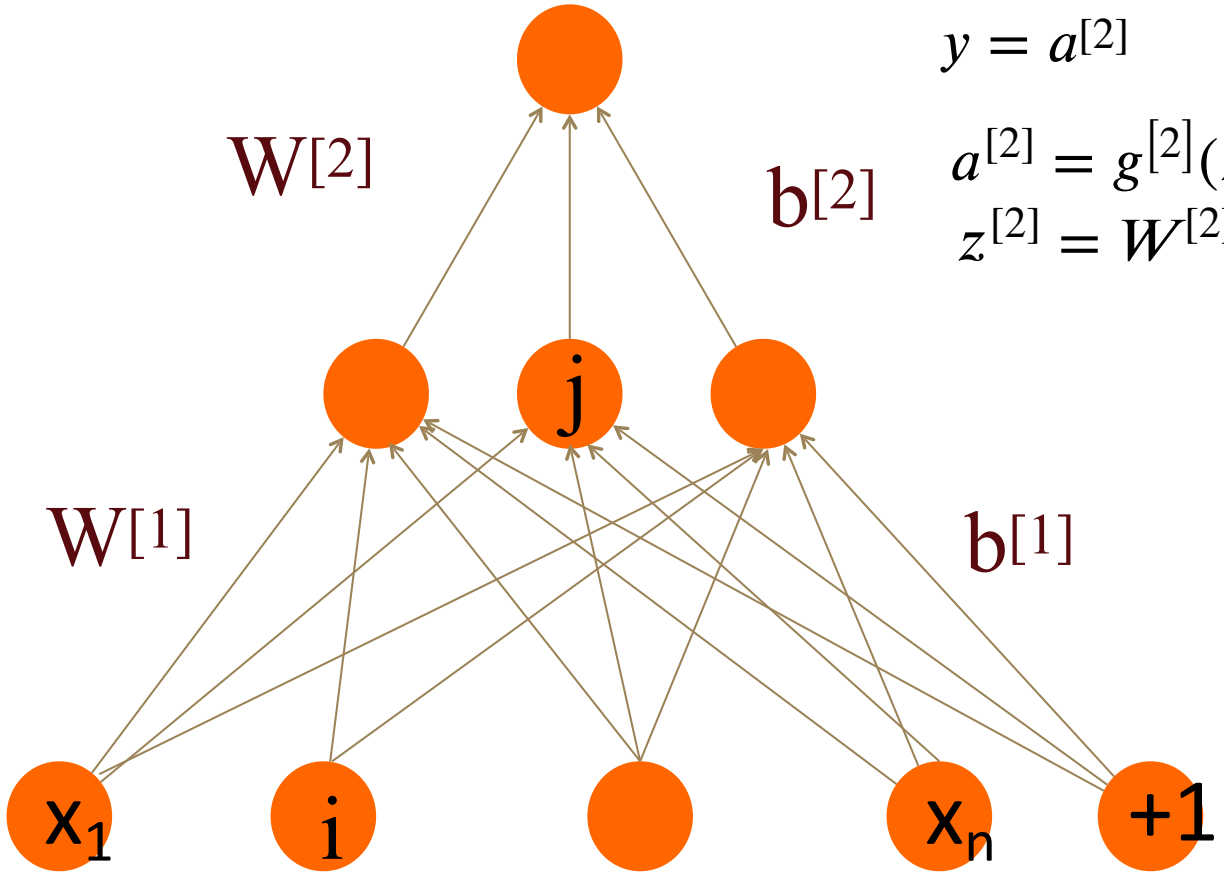
Two-Layer Network with scalar output



Two-Layer Network with softmax output



Multi-layer Notation



$y = a^{[2]}$ sigmoid or softmax

$a^{[2]} = g^{[2]}(z^{[2]})$

$z^{[2]} = W^{[2]}a^{[1]} + b^{[2]}$

$a^{[1]} = g^{[1]}(z^{[1]})$ ReLU

$z^{[1]} = W^{[1]}a^{[0]} + b^{[1]}$

$a^{[0]}$

Replacing the bias unit

Let's switch to a notation without the bias unit (a notational change):

1. Add a dummy node $a_0=1$ to each layer
2. Its weight w_0 will be the bias
3. So input layer $a^{[0]}_0=1$,
   - And $a^{[1]}_0=1$, $a^{[2]}_0=1$,…

Replacing the bias unit

w_0 = weight applied to bias feature

Instead of:

$$x = x_1, x_2, \ldots, x_{n0}$$

$$\mathbf{h} = \sigma(\mathbf{W}\mathbf{x} + \mathbf{b})$$

$$\mathbf{h}_j = \sigma\left(\sum_{i=1}^{n_0} \mathbf{W}_{ji}\,\mathbf{x}_i + \mathbf{b}_j\right)$$

We'll do this:

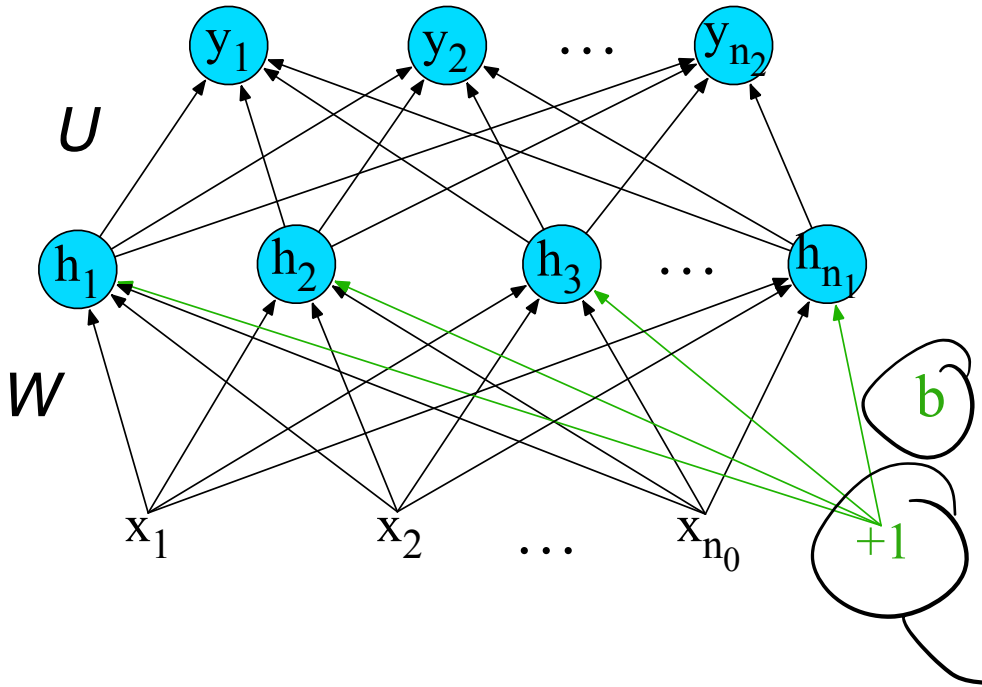
$$x = x_0, x_1 \cdots x_n$$

$$h = \sigma(\underline{W}\underline{x})$$

$$h_j = \sigma\left(\sum_{i=0}^{no} W_{ji}\, x_i\right)$$

Replacing the bias unit

Instead of:



We'll do this:

