Machine Learning

Please evaluate this course on Axess!

Your comments really do make a difference.

Announcements

- Assignment 6 due right now.
  - Due Friday at 3:15PM with one late day.
  - Due Monday at 3:15PM with two late days.
- Assignment 7 (**FacePamphlet**) out now, due next **Thursday, March 21** at 3:15PM
  - **No late submissions accepted** (sorry about that – this is university policy).
  - **This is a hard deadline: anything received at 3:15:01PM or later will not be accepted.**
- We will be holding extra LaIR hours this Sunday through Wednesday from 7PM – 11PM each night.

FacePamphlet Demo

Perceptron Learning

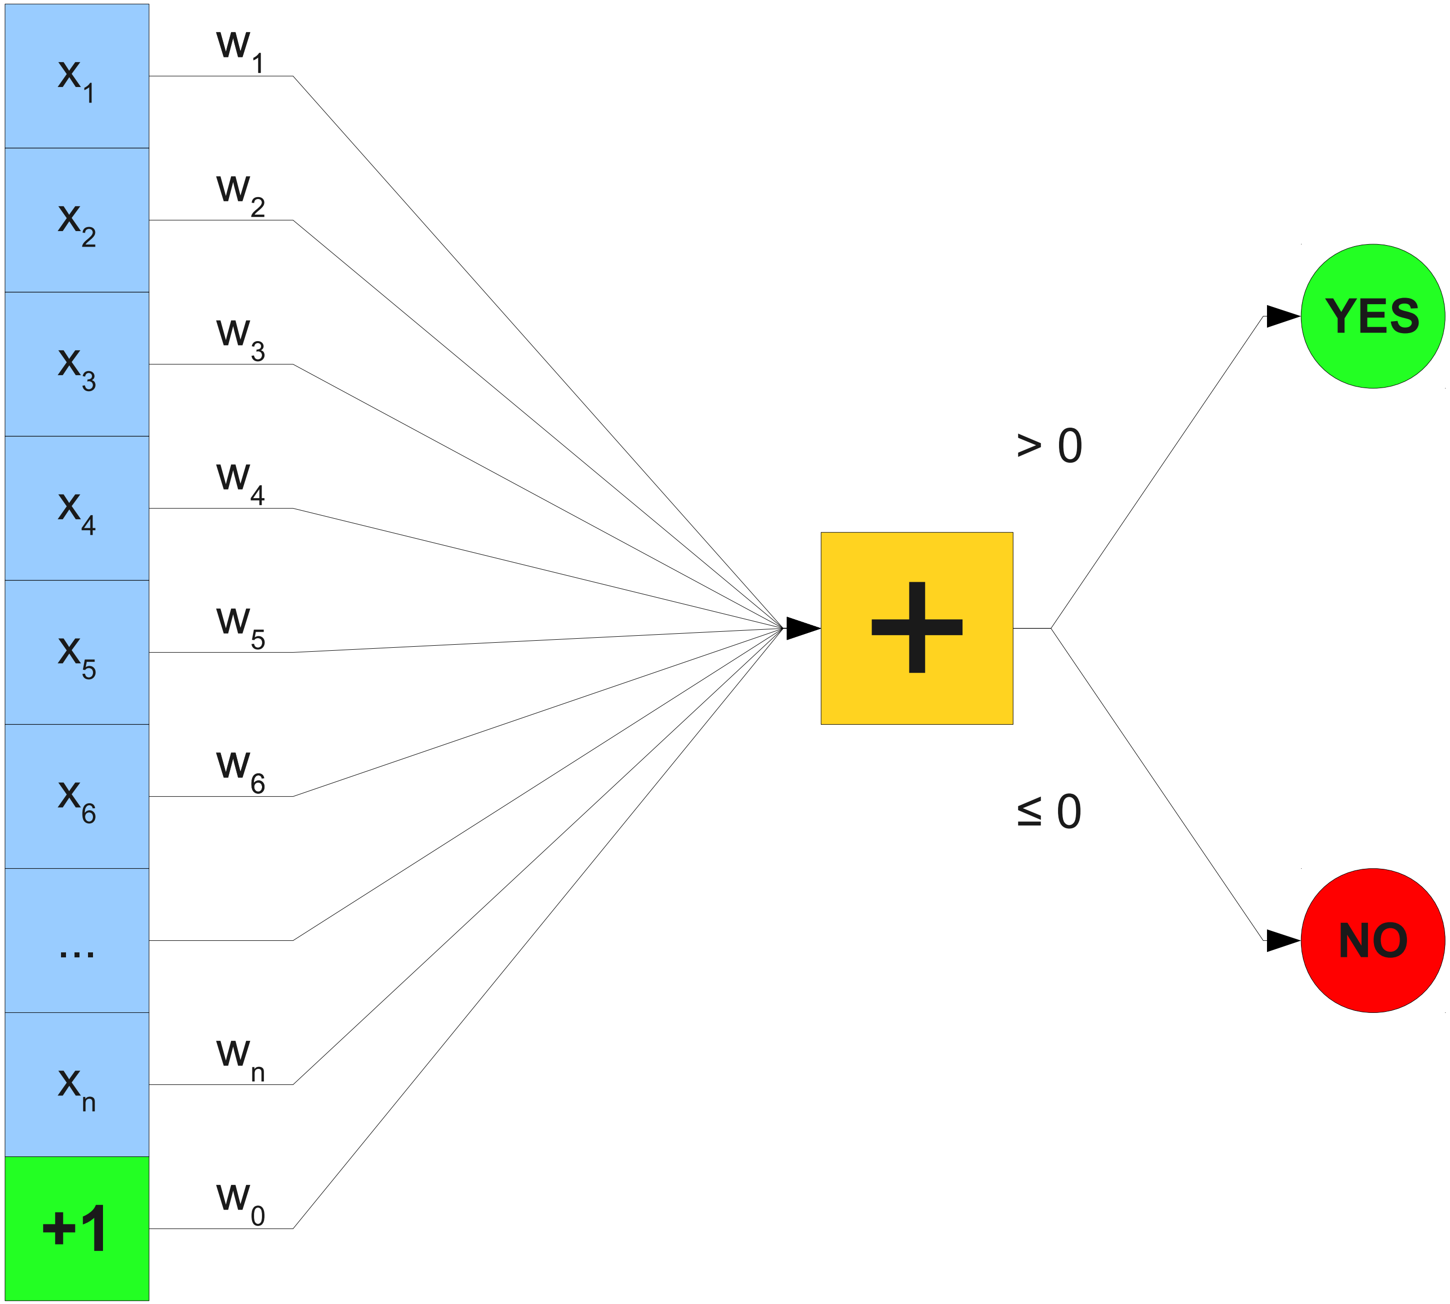

x1
x2
x3
x4
x5
x6
...
xn
+1
w1
w2
w3
w4
w5
w6
wn
w0
+
> 0
≤ 0
YES
NO

How do we choose good values for w_0 … w_n?

One Approach

- **Train** the perceptron on valid data.
- For each data point:
  - Ask the perceptron what it thinks.
  - If correct, do nothing.
  - Otherwise, nudge w_0 … w_n in the right direction.
- Repeat until number of errors is “small enough.”
- Question: What kind of mistakes can we make?

False Positive

False Positive

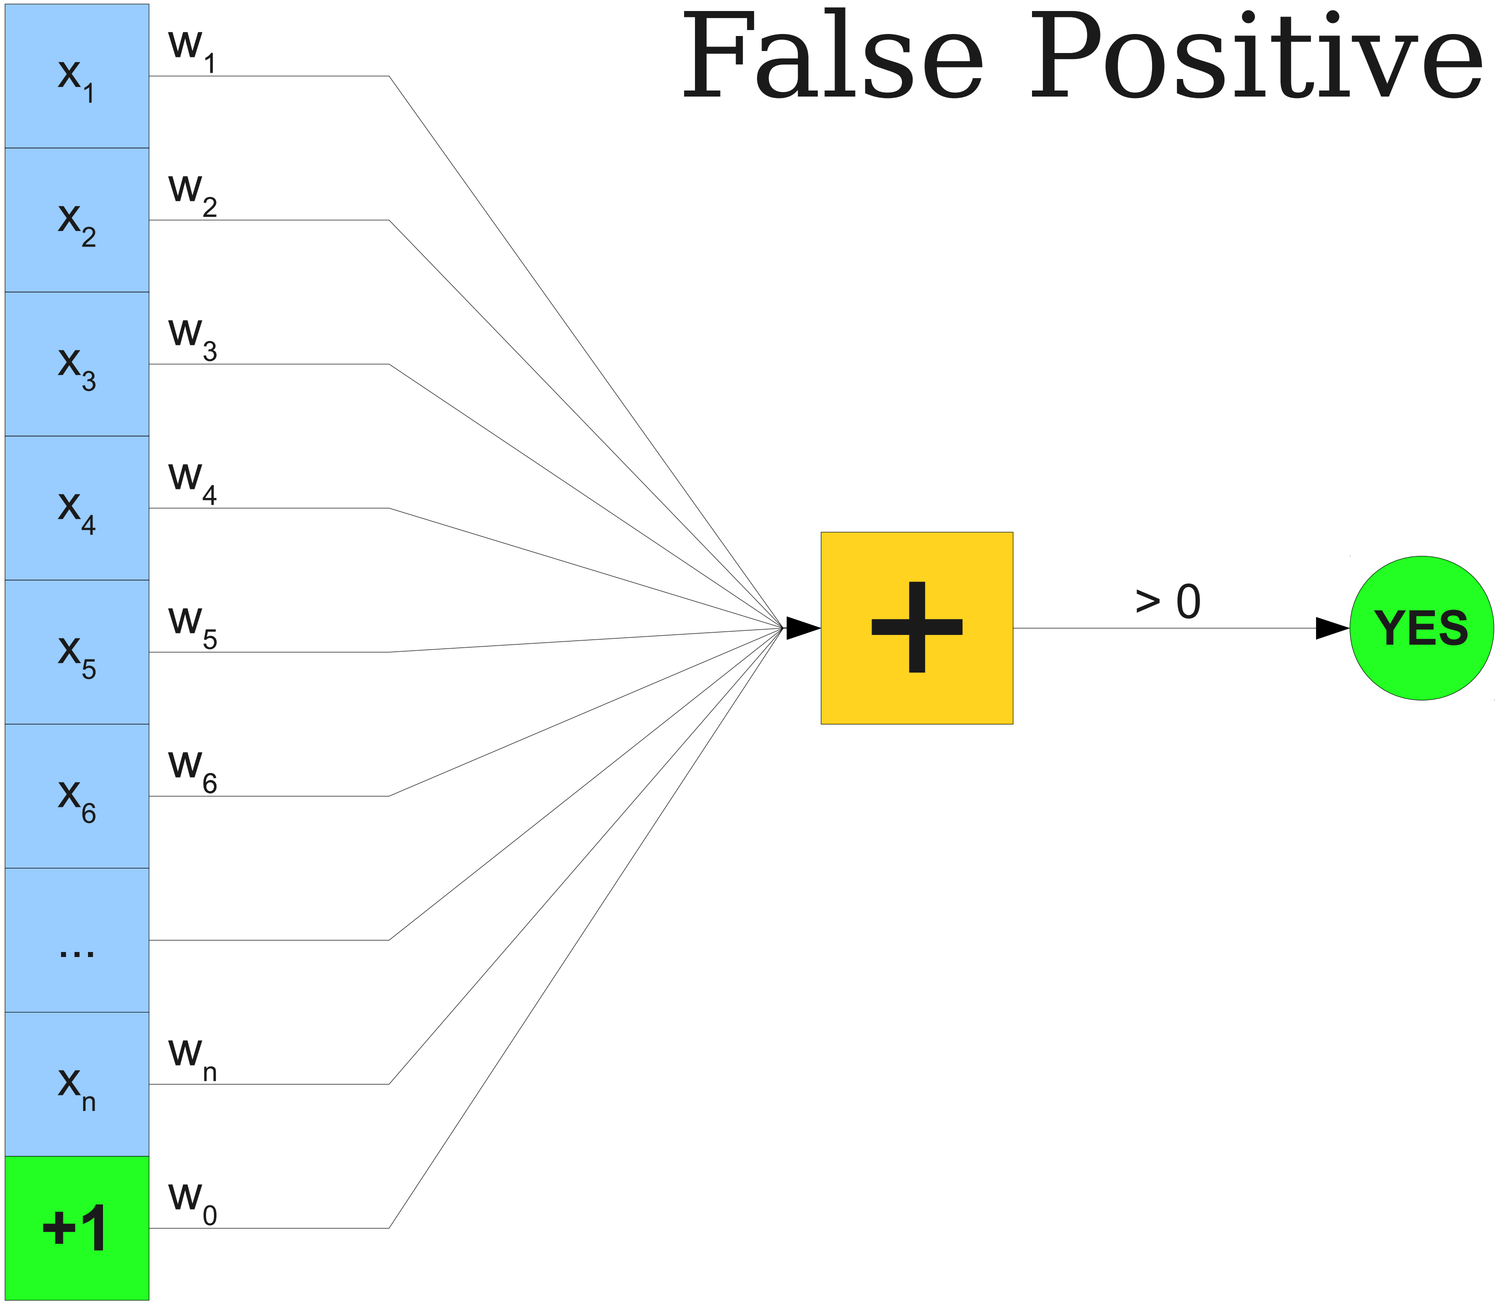

False Positive

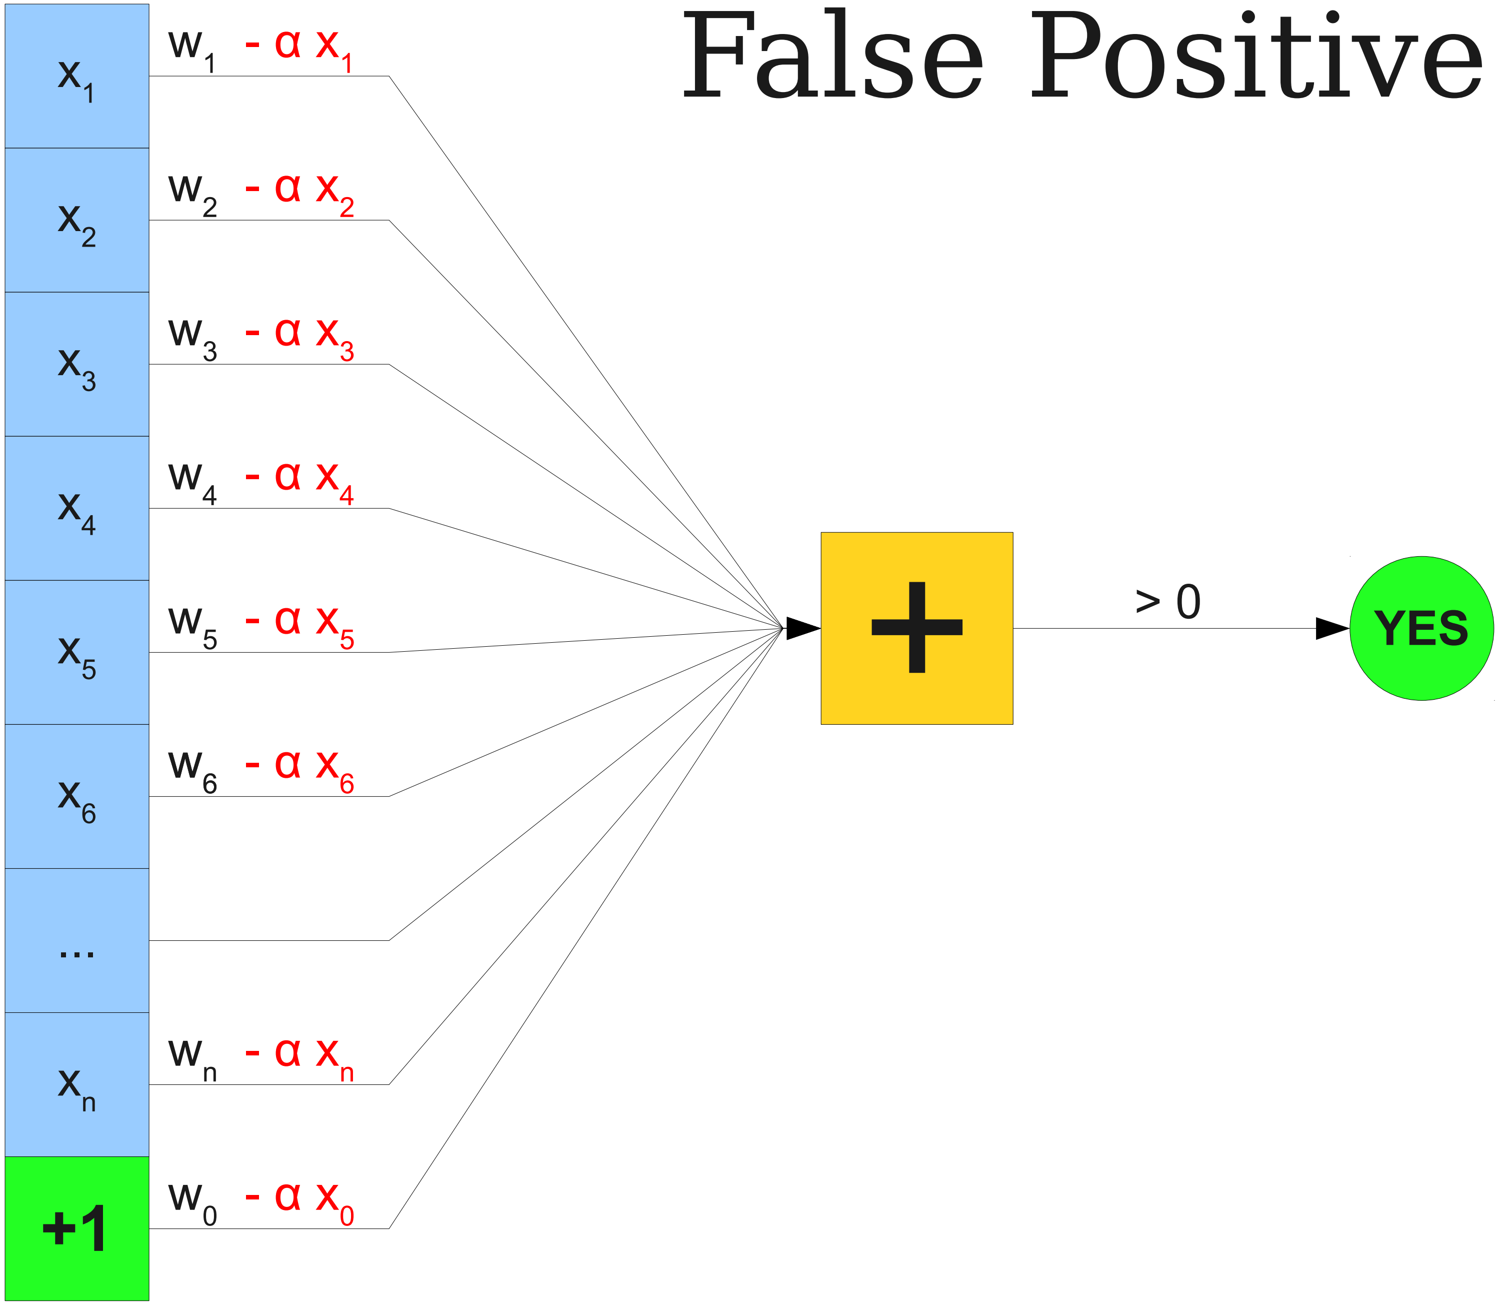

False Negative

False Negative

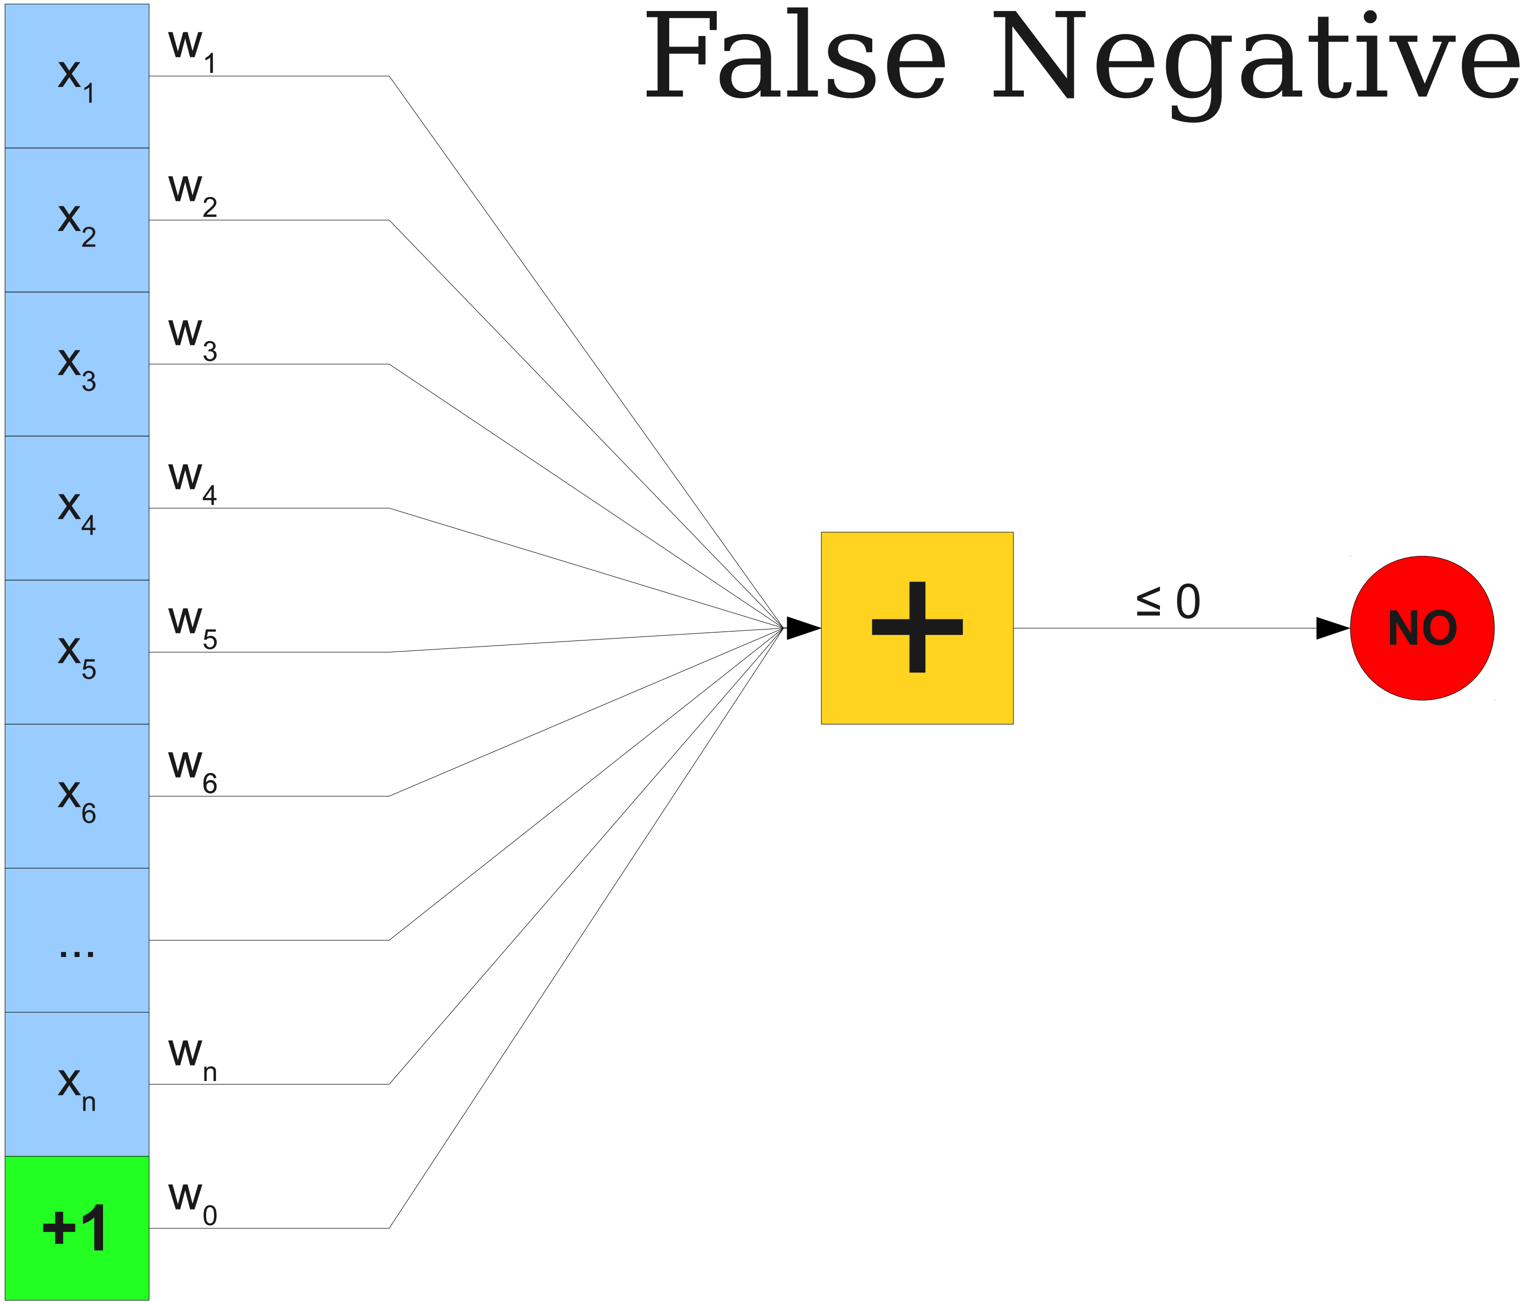

False Negative

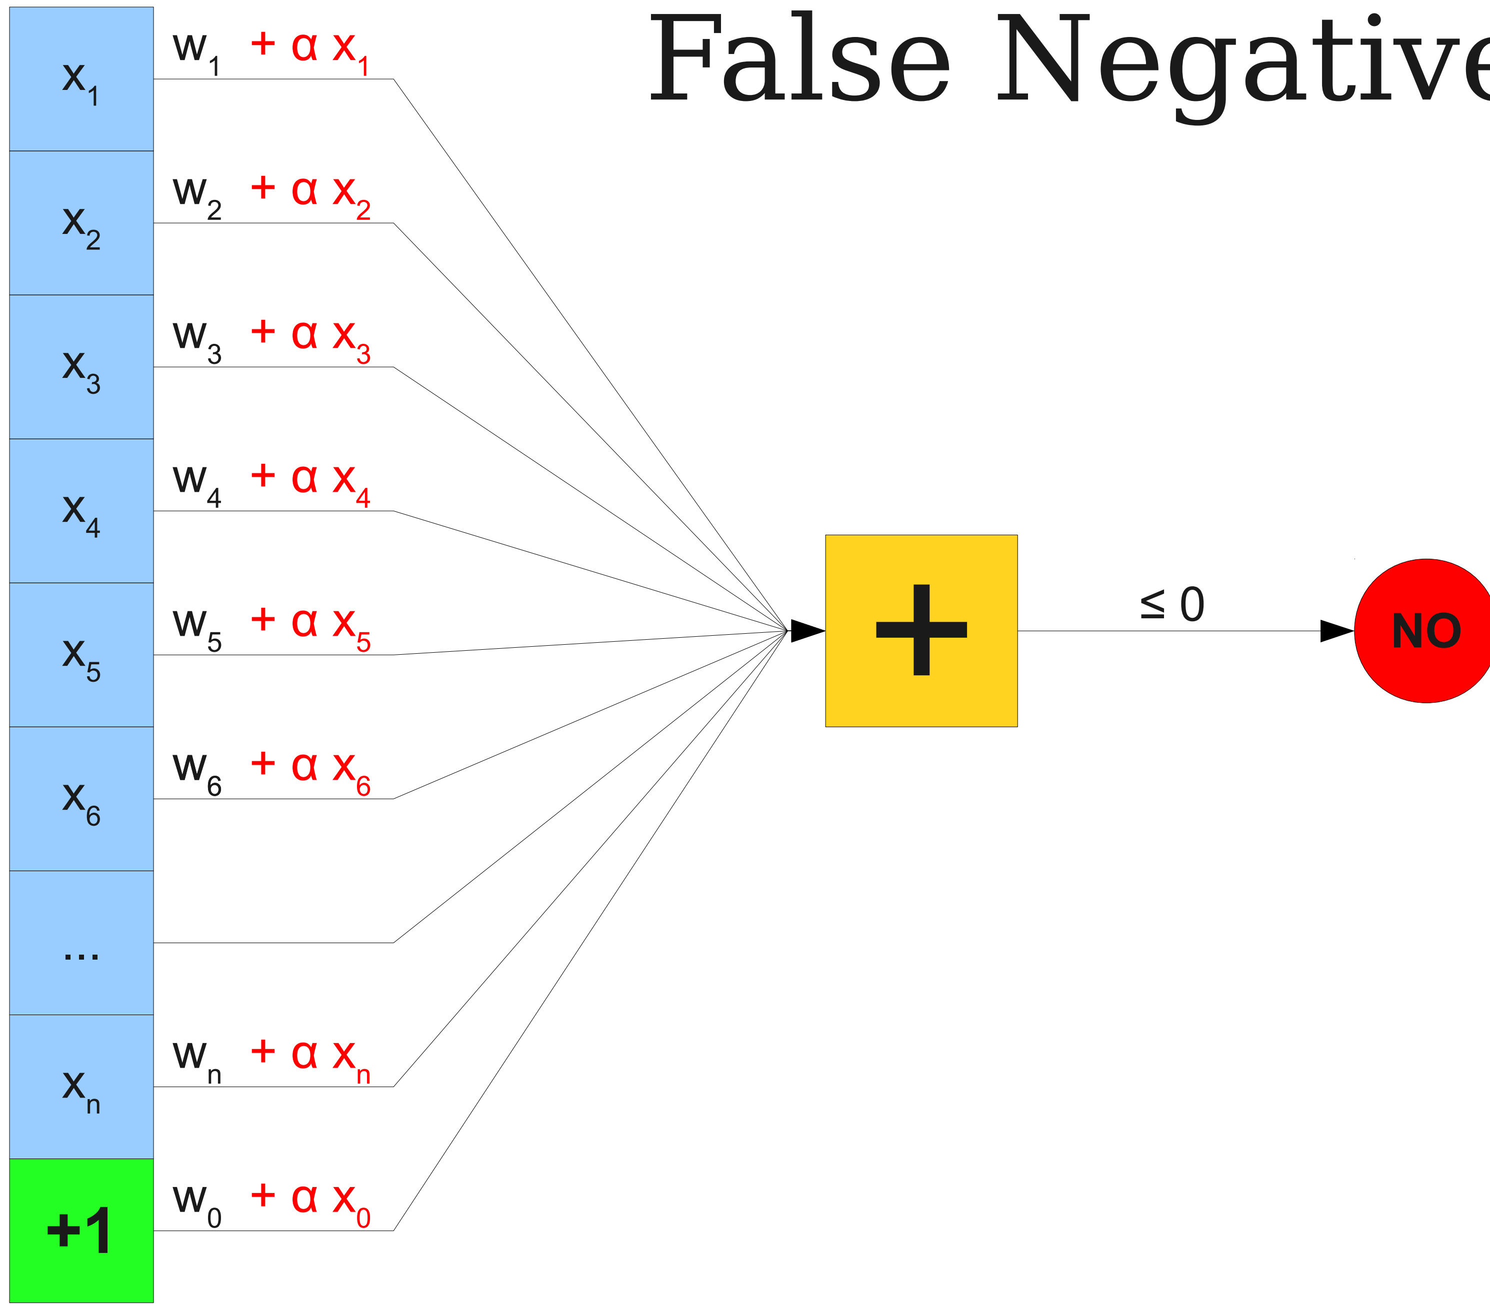

A Cute Math Trick

- For false positives, set $w_i = w_i - \alpha\, x_i$.
- For false negatives, set $w_i = w_i + \alpha\, x_i$.
- For correct answers, set $w_i = w_i$.
- Let “YES” be 1 and “NO” be 0.
- Consider the difference between **actual answer** and **perceptron guess**:
  - False positive: Actually NO, we say YES. Difference is -1.
  - False negative: Actually YES, we say NO. Difference is +1.
  - Correct answer: Both YES or both NO. Difference is 0.
- General update rule: $\boldsymbol{w_i = w_i + \alpha\ (\text{real - guess})\ x_i}$.

Perceptron Learning Algorithm

- Start with a random guess of each w_i.
- Repeat until perceptron is sufficiently accurate:
  - Choose a training example $(x_0, x_1, \ldots, x_n)$.
  - Let **real** be the real answer, **guess** be the perceptron's guess.
  - For each i, set $\boldsymbol{w_i = w_i + \alpha\ (\text{real - guess})\ x_i}$
- Note: Use **batching** in practice.
  - Update everything all at once.

Application: Handwriting Analysis

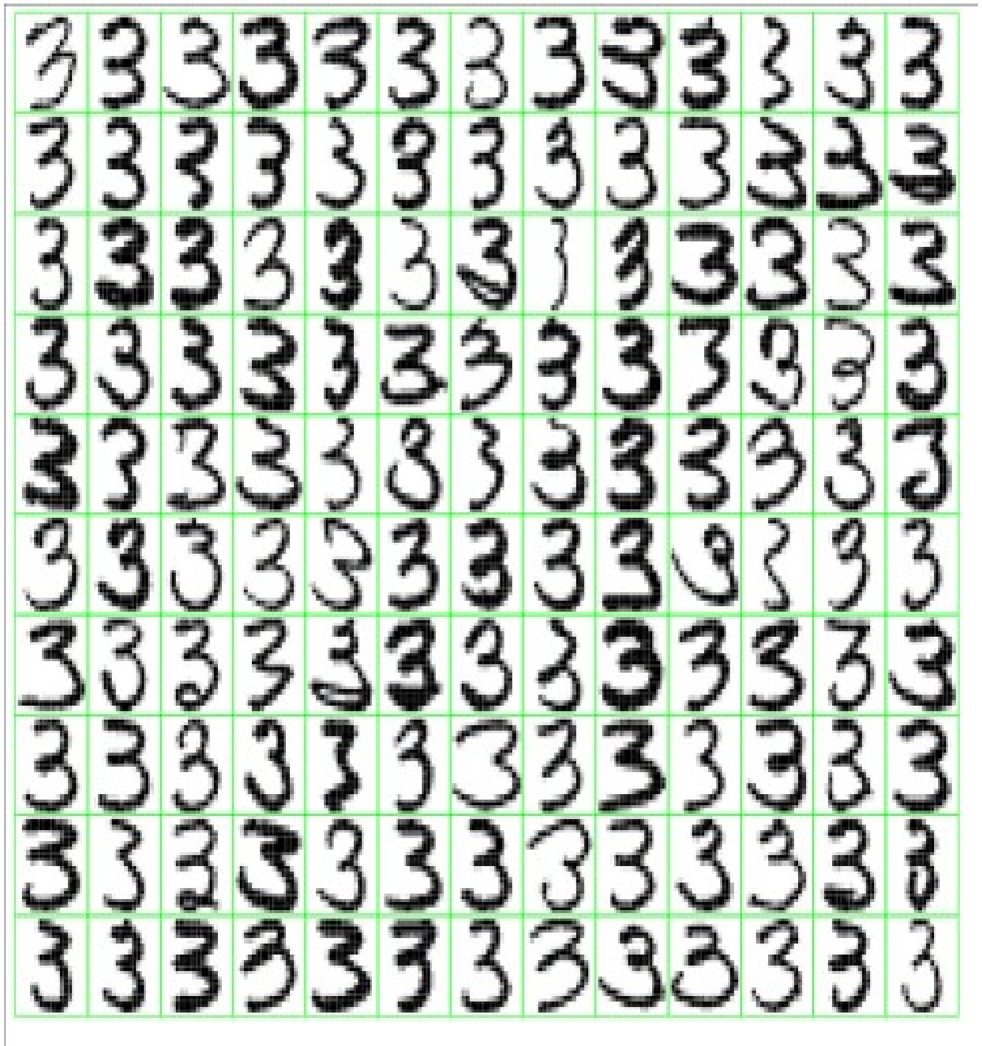

- Train a computer to recognize handwritten numbers 0 – 9.
- Large training and test set available (MNIST Handwritten Digit Database)

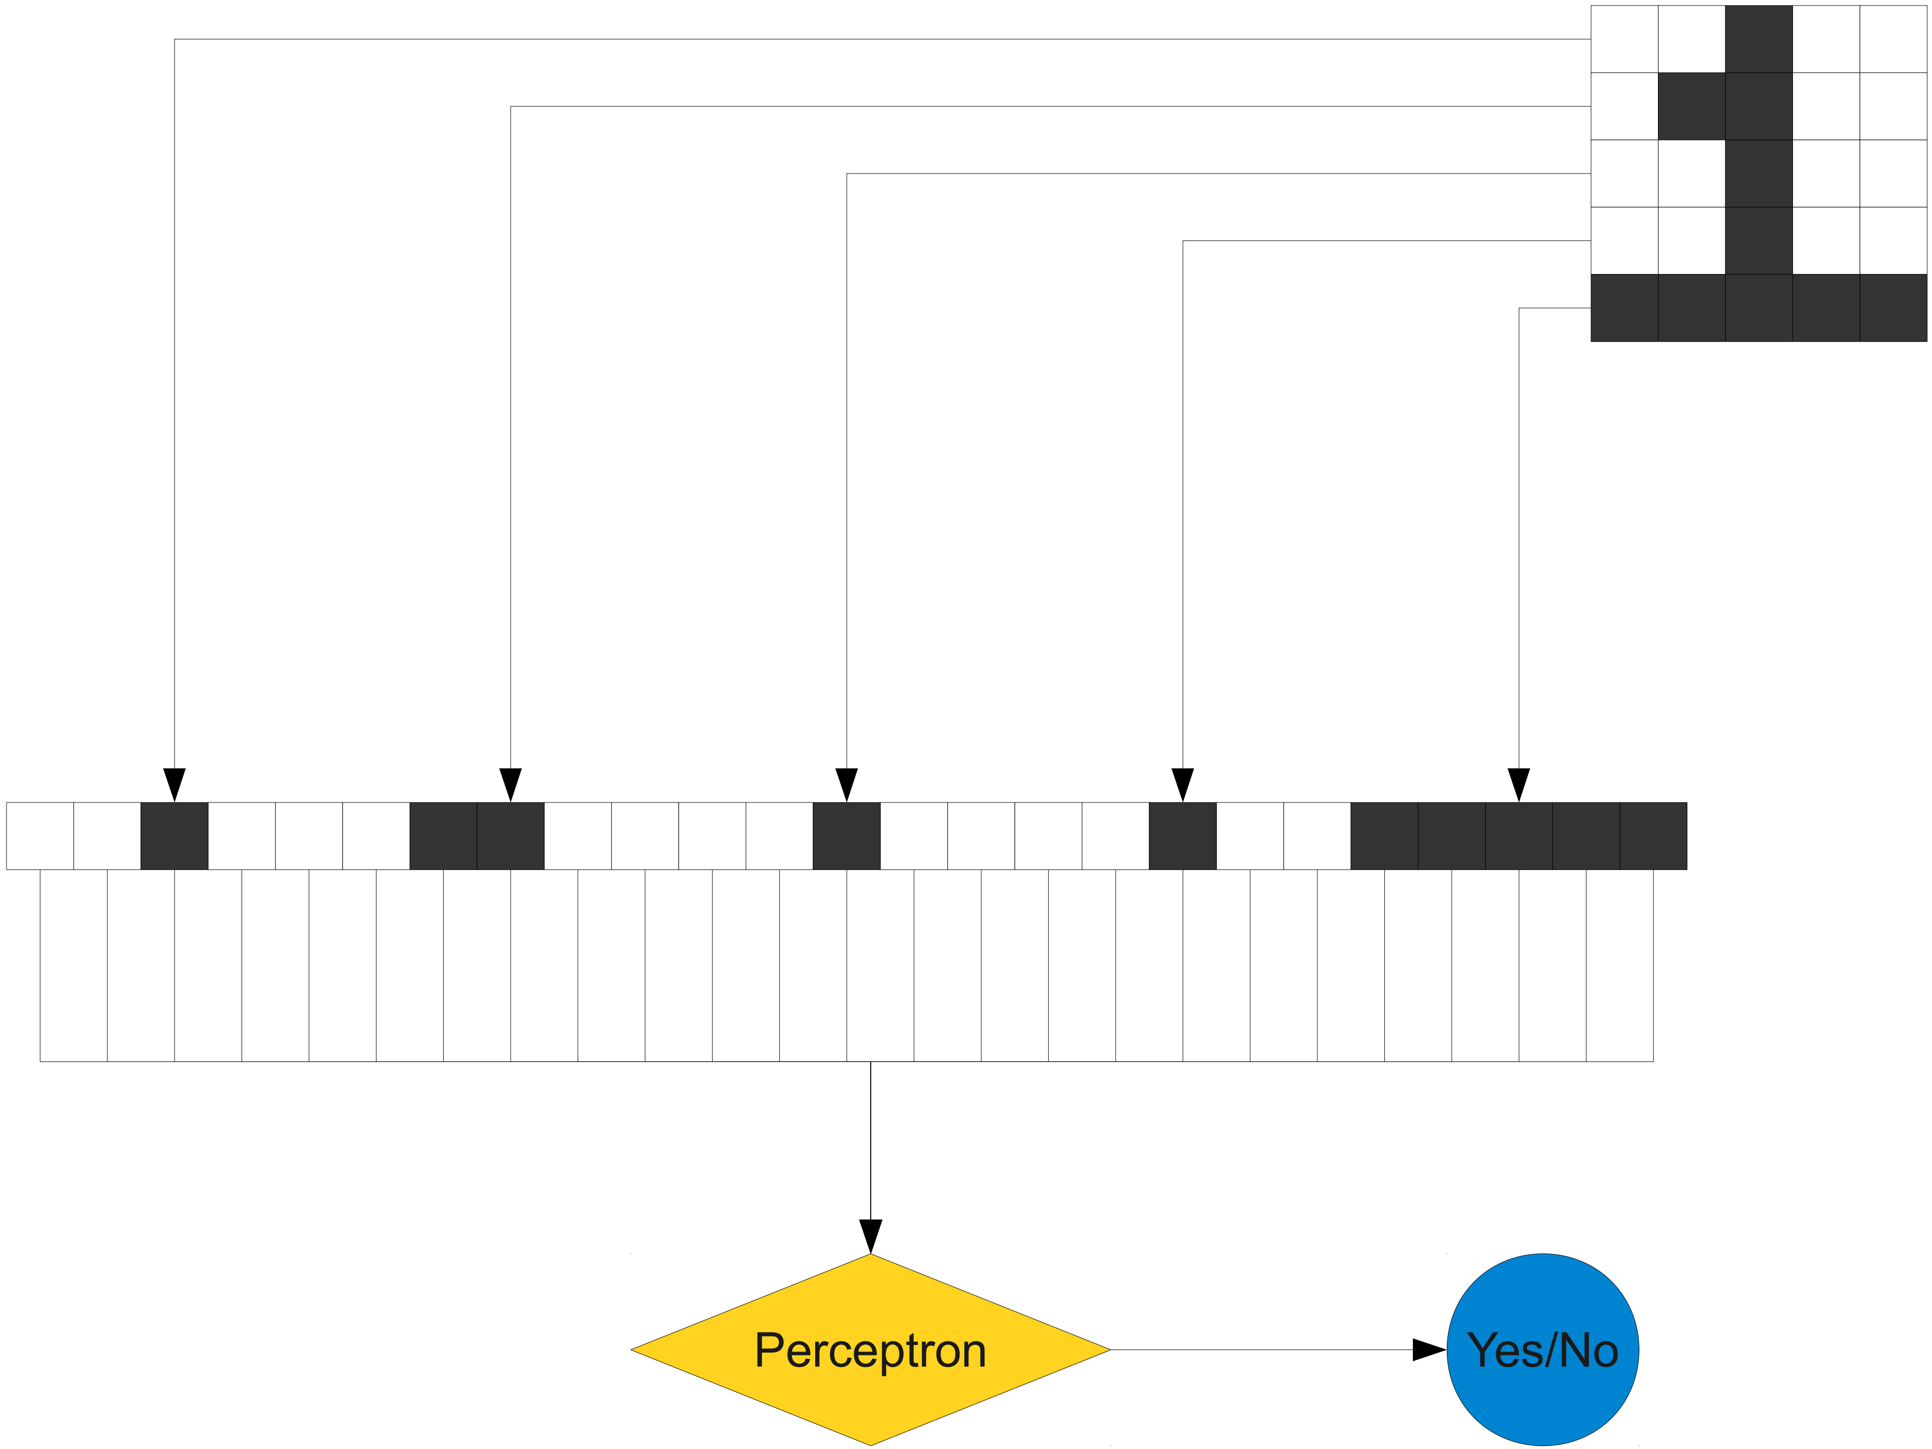
Perceptron
Yes/No

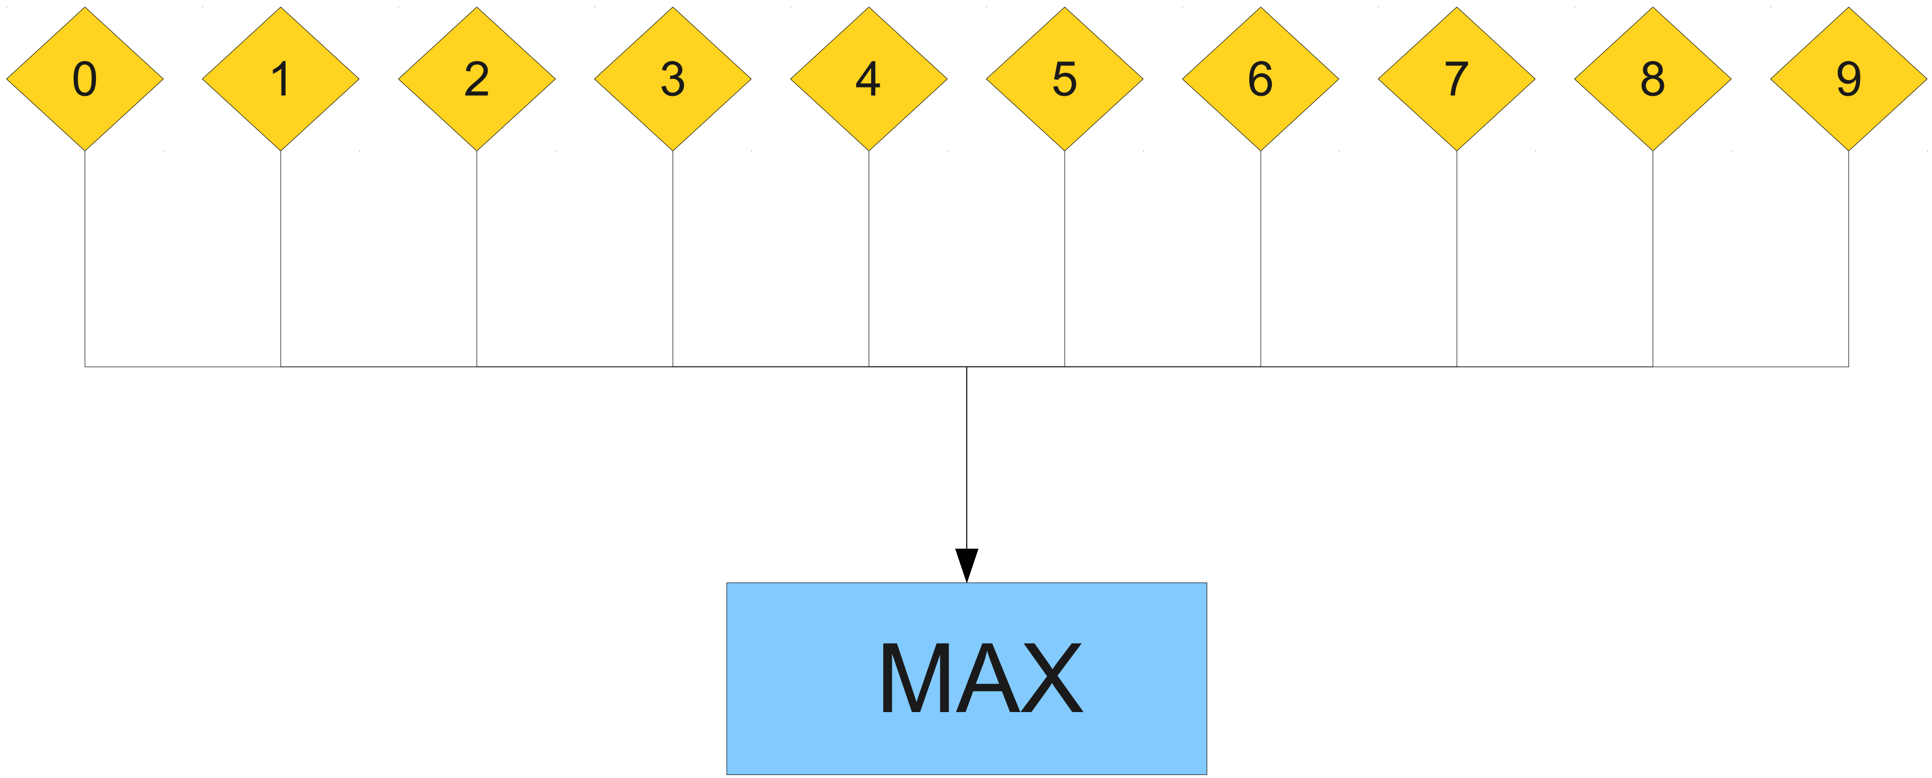
0
1
2
3
4
5
6
7
8
9
MAX

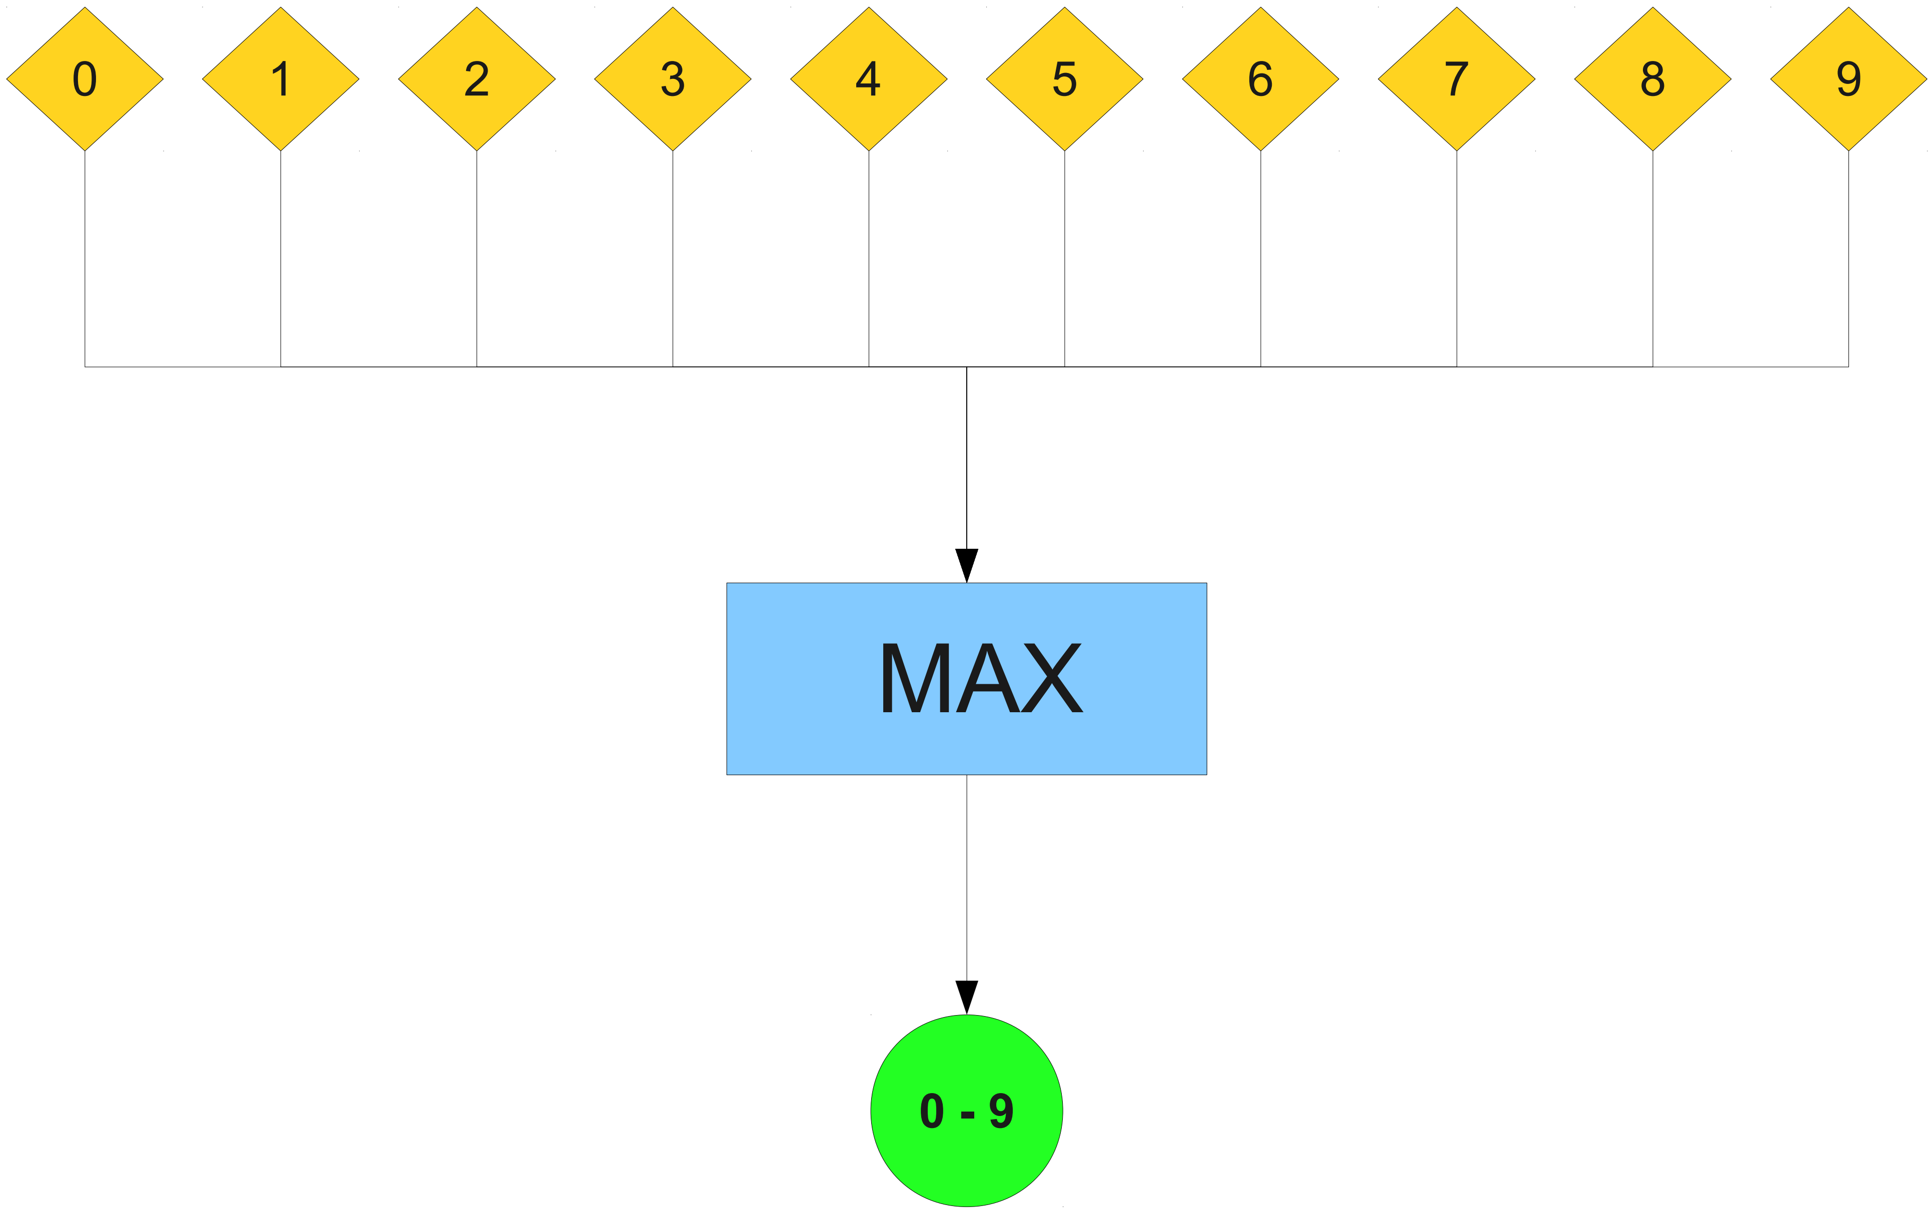
0
1
2
3
4
5
6
7
8
9
MAX
0 - 9

Combining Perceptrons

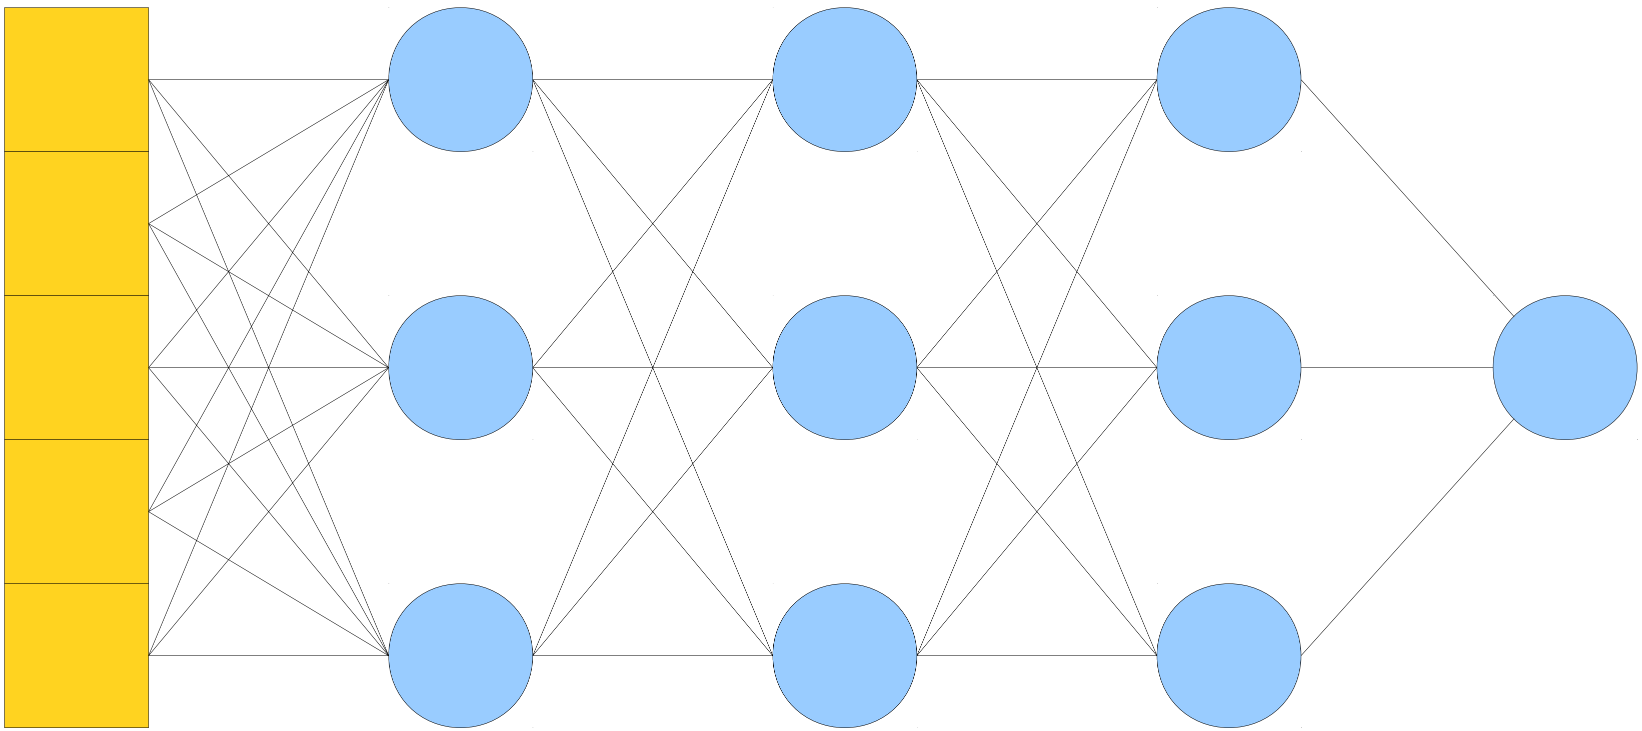

This is called a **neural network**.

Machine Learning

- Interesting in machine learning? Take CS109 and CS229!
- Many beautiful algorithms:
  - Naive Bayes classifiers (used in spam filtering).
  - Decision trees (used in hospitals for diagnostics).
  - Bayesian networks (used in cancer research to learn what causes tumors).
  - Restricted Boltzmann machines (used to learn what cats look like).
    - http://www.nytimes.com/2012/06/26/technology/in-a-big-network-of-computers-evidence-of-machine-learning.html?pagewanted=all

Next Time

- **Where to Go from Here**
  - What comes next in CS?
  - What can you do with what you've learned?