Computers vs. brains

- The design of modern computers is partly based on how brains were thought to work (e.g. von Neumann, 1940s)
- While our understanding of brains is influenced by our model of computation
- The theoretical computing power of computers and brains is in a sense the same (McCulloch and Pitts, 1940s)
- But actual brains and actual computers work very differently!
- Computers:
    - One very fast central processor (CPU; cf. Turing machines)
    - Memory separate from processing
- Brains:
    - Many units computing in parallel
    - Memory distributed throughout the entire network
- E.g. perception occurs in <100 cycles vs. millions in a CPU

Pandemonium (Selfridge, 1959)

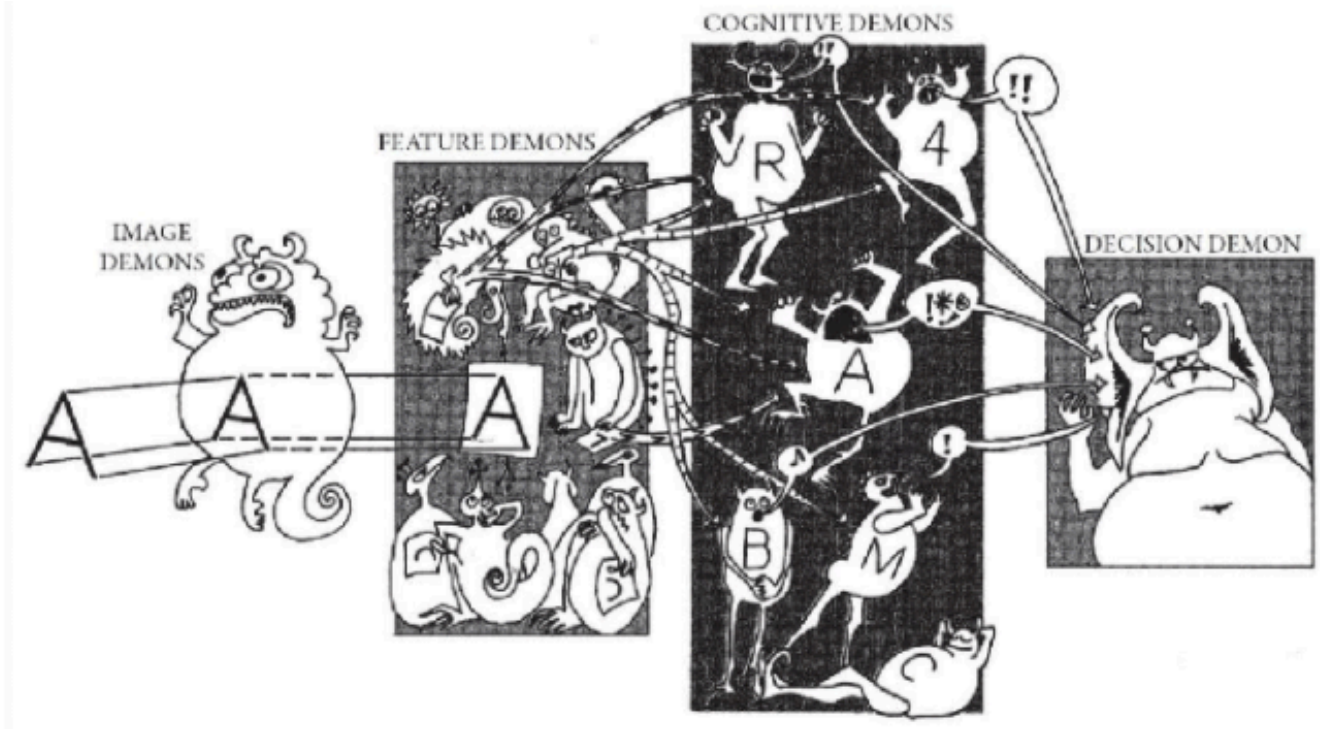

A network of "demons" all talking to each other at the same time

Hebbian learning

- The key question is how the brain learns, i.e. how it modifies the weights along the connections.
- Donald Hebb (1949) proposed that the brain learns by strengthening connections between co-occurring events

  - A neural version of Locke's idea of strengthening associations between experiences

Classification bound models

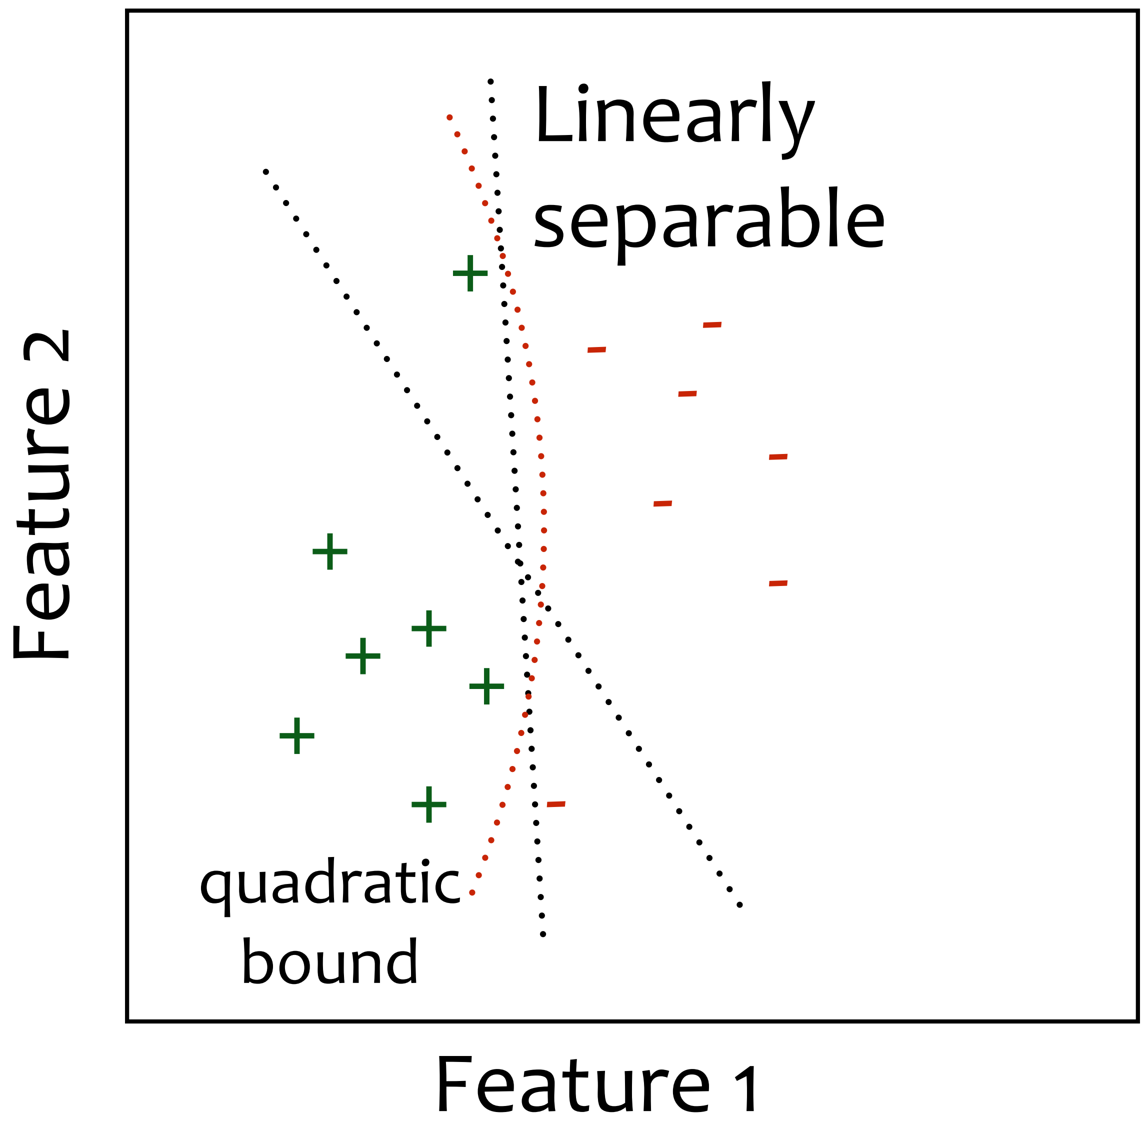

- Learn a classification boundary between categories
- Classify by applying bound
- Linear bounds are particularly easy
- Hypothesis: humans can handle quadratic bounds (linear bounds are a special case)

Perceptrons

- A perceptron is a (generalized) neuron

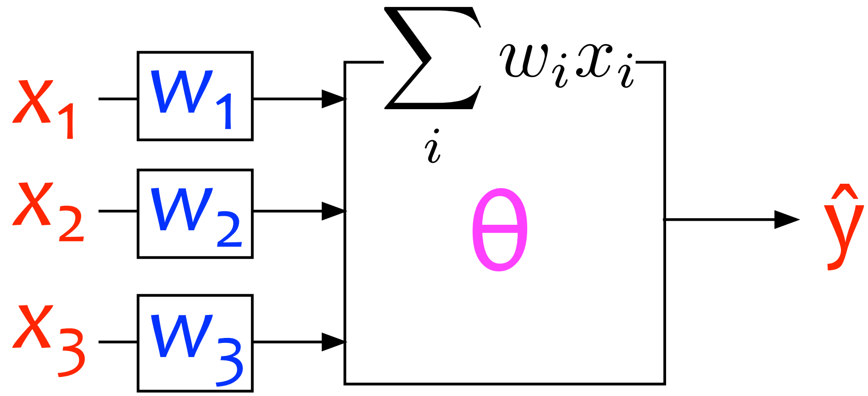

- It has

  set of input wires $x_i = x_1 \ldots x_n$ (input space X, aka "retina.")

  each of which has a weight w_i,

  a threshold θ,

  and a single output ŷ.

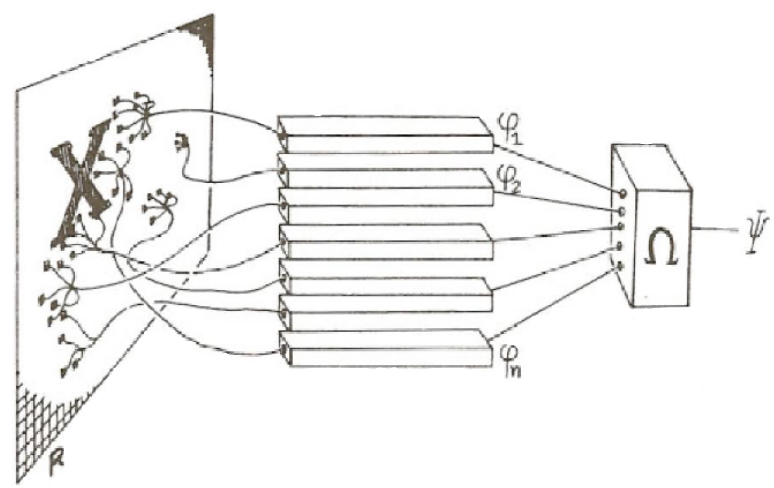

Perceptrons as classifiers

- The weighted sum $y = w_1 x_1 + w_2 x_2 + \ldots =$

  $y = \sum_i w_i x_i$

  defines a line

  (or plane or hyperplane...)

  in the input space X.

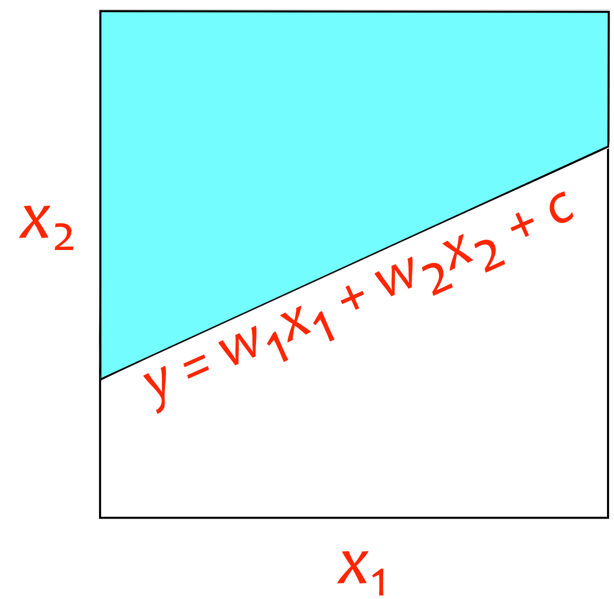

- The inequality

  $\sum_i w_i x_i > \theta$

  defines a region in the input space X.

- So the perceptron classifies the input

Linear separability

- Rosenblatt (1962) demonstrated a procedure (the perceptron learning rule) that could adjust the weights w_i to find the best classifier of the input data.

- But Minsky & Papert (1969) realized that this meant that perceptrons could only classify linearly separable problems

  ...but some very important problems are not linearly separable

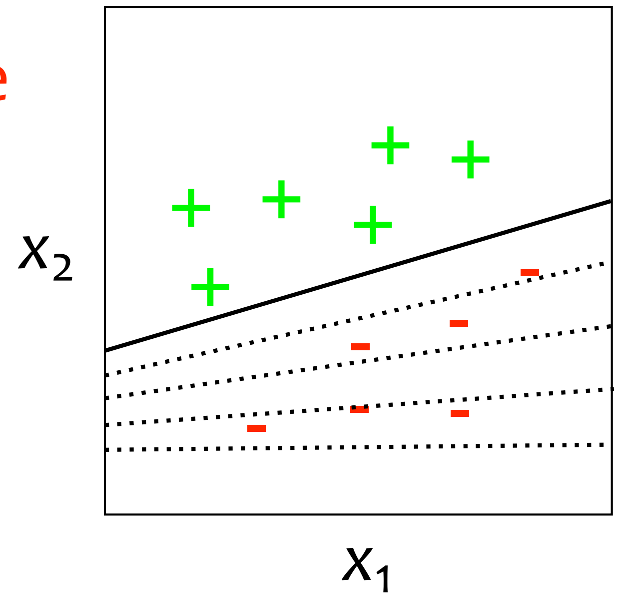

- This killed research in artificial neural networks for about 15 years

Limits on perceptrons

"XOR"

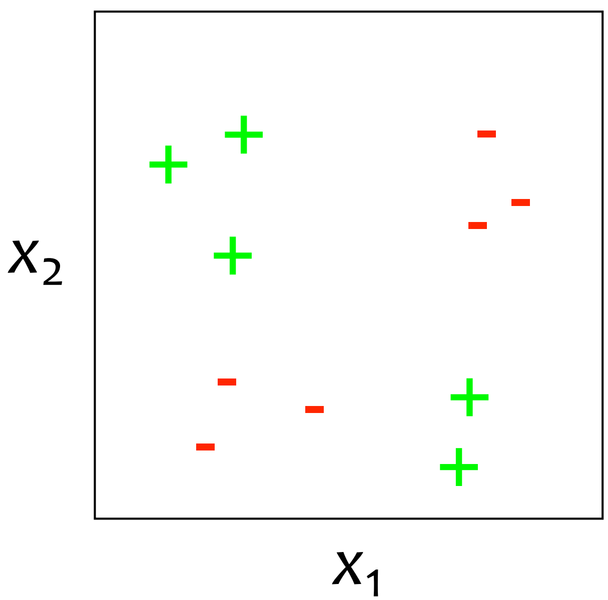

not linearly separable

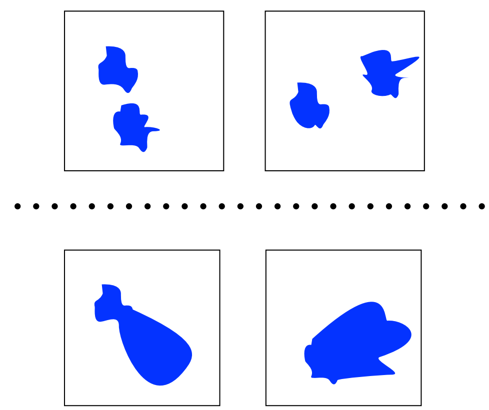

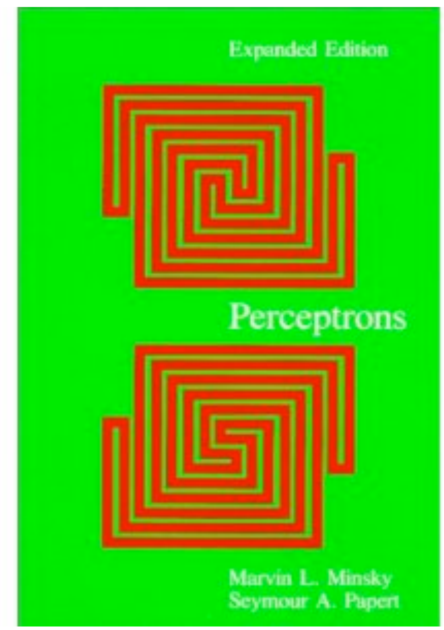

connected vs. unconnected

can't be classified linearly if receptive fields are local

= can't be "parallelized"

Classical machine learning

- Machine Learning is the field of AI in which computers are programmed to learn from examples
- During the 1970s and 1980s, much of AI involved the application of complex rule systems, often "subjectively" programmed

  - "Good old fashioned AI" (GOFAI)

  - Expert systems — approximate human experts as combinations of conditional rules and procedures

- …including in Machine Learning, where learning involved inducing the right set of if-then statements

Learning by inducing rules

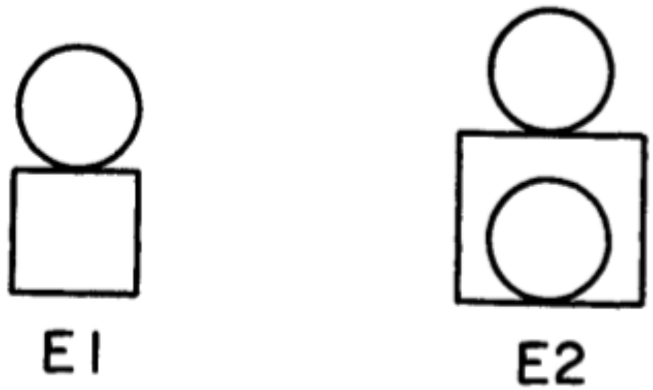

Positive examples

- First, find the right description language
  - E1: `ontop(circle1,square)`
  - E2: `ontop(circle1,square);inside(circle2,square)`
- Drop condition `inside(circle2,square)`
- Induced rule:
  - `IF ontop(circle,square) THEN ispositive`