CSE 473: Artificial Intelligence

Multilayer Perceptrons

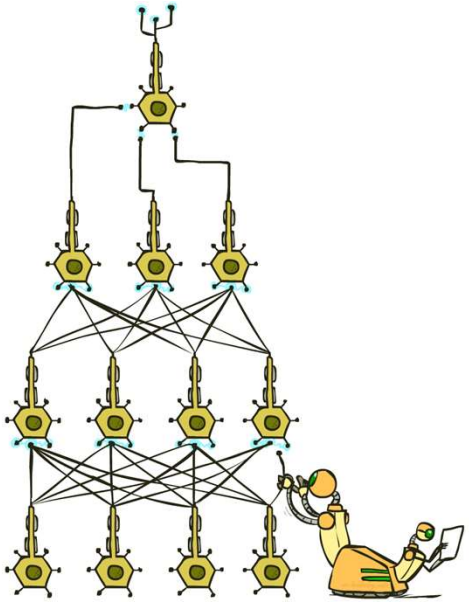

Steve Tanimoto --- University of Washington

How Many Computers to Identify a Cat?

How Many Computers to Identify a Cat? 16,000

An image of a cat that a neural network taught itself to recognize.

By JOHN MARKOFF
Published: June 25, 2012

MOUNTAIN VIEW, Calif. — Inside Google's secretive X laboratory, known for inventing self-driving cars and augmented reality glasses, a small group of researchers began working several years ago on a simulation of the human brain.

There Google scientists created one of the largest neural networks for machine learning by connecting 16,000 computer processors, which they turned loose on the Internet to learn on its own.

Presented with 10 million digital images found in YouTube videos, what did Google's brain do? What millions of humans do with YouTube: looked for cats.

"Google Brain"

[Le, Ng, Dean, et al, 2012]

Perceptron

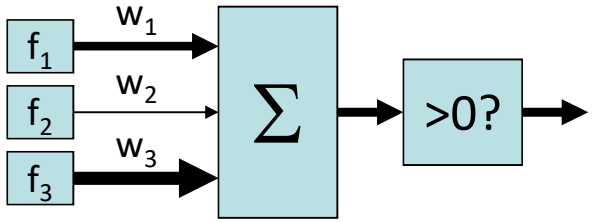

Two-Layer Neural Network

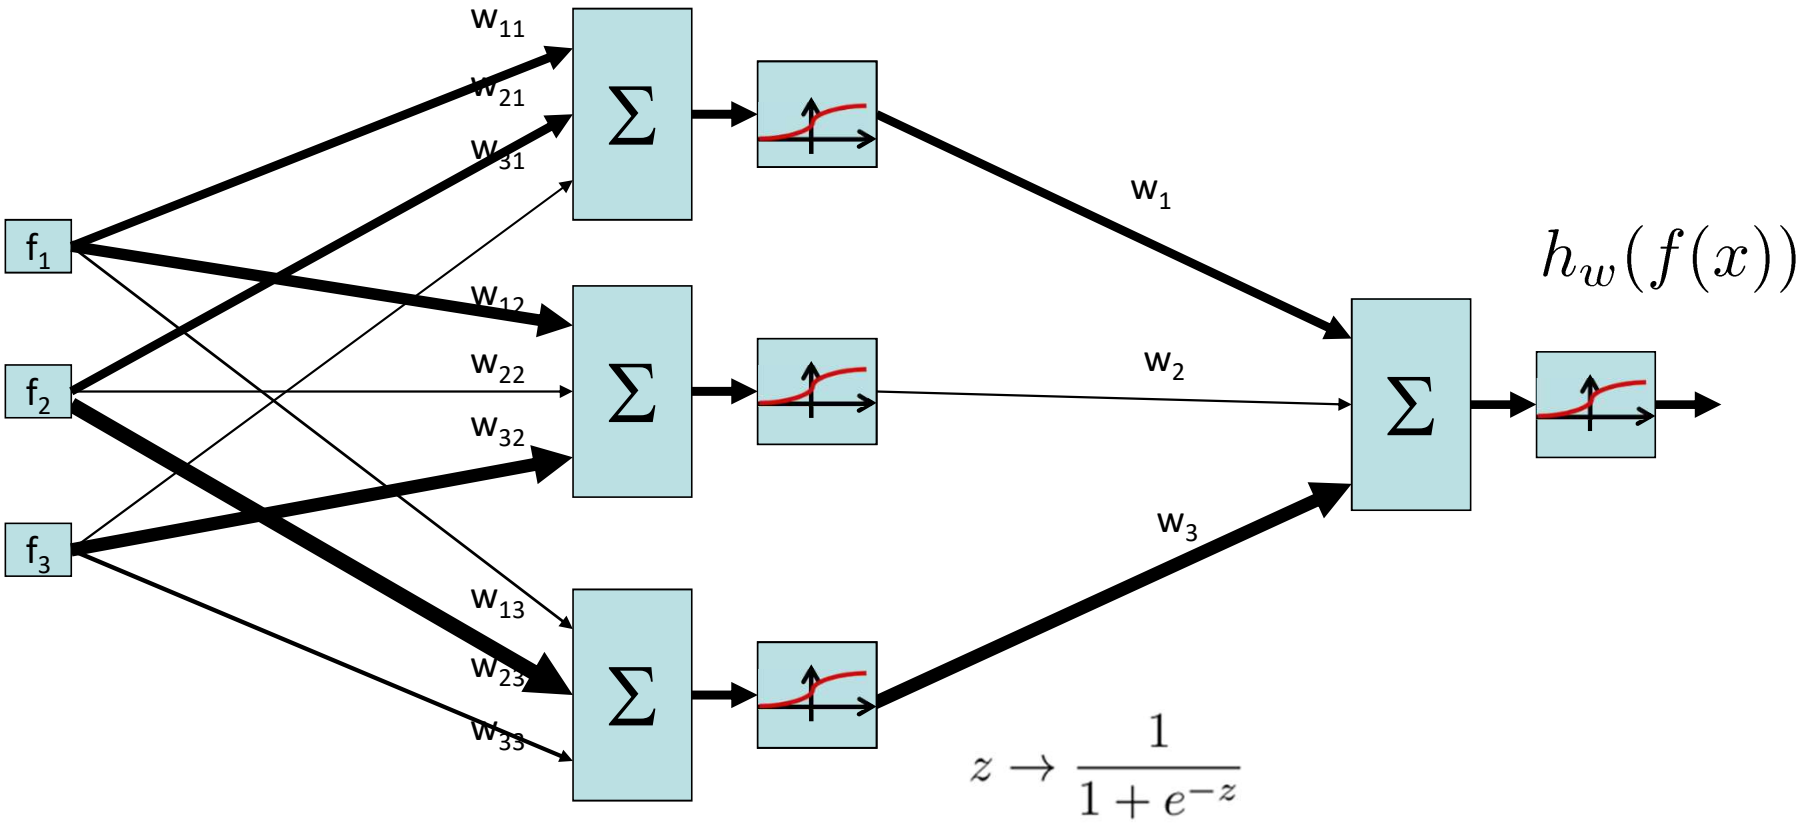

N-Layer Neural Network

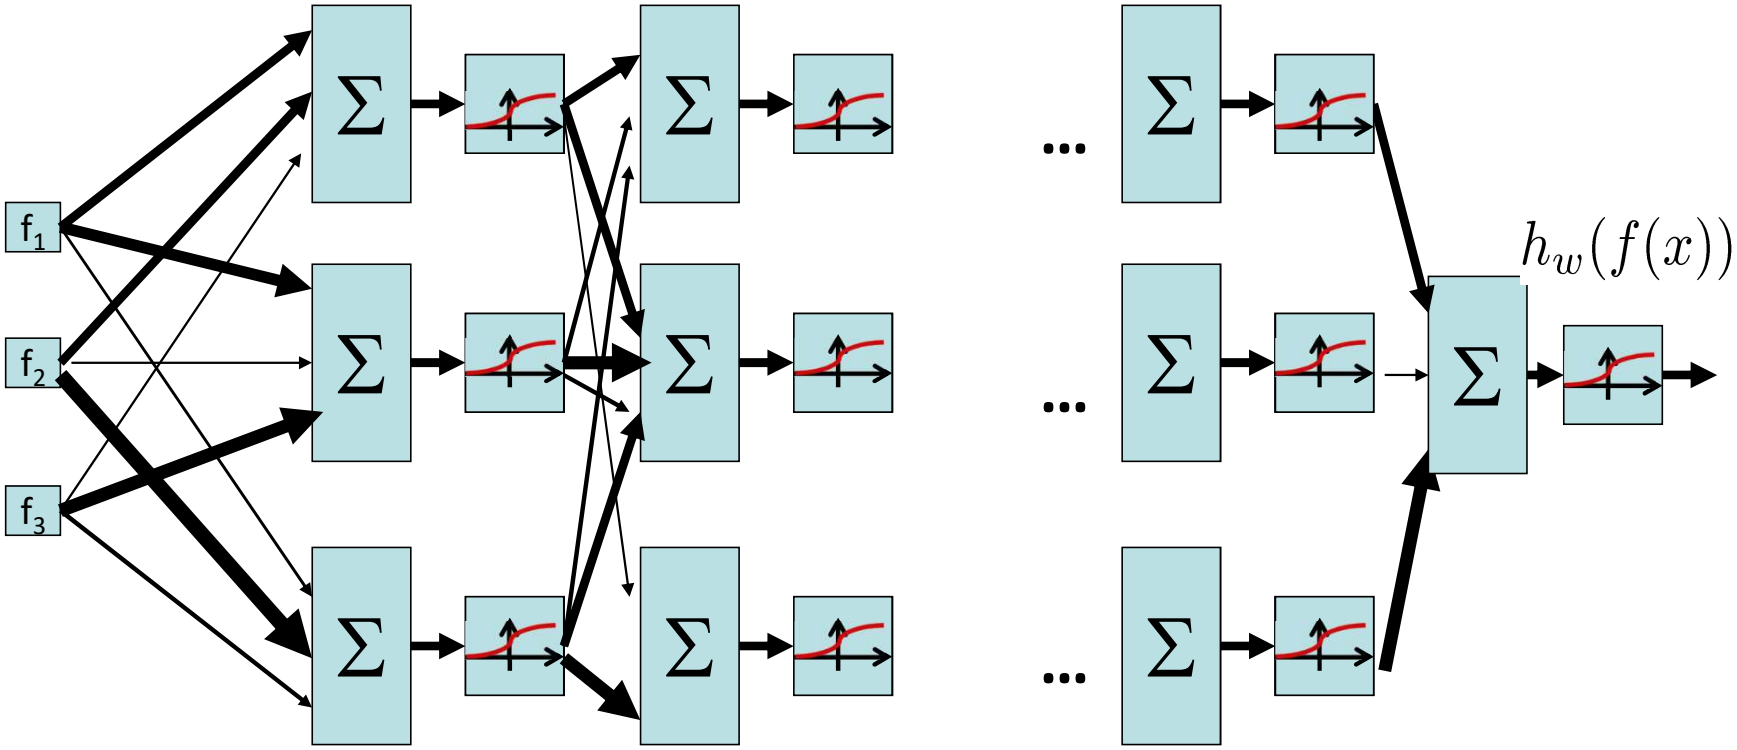

Hill Climbing

- Simple, general idea:
  - Start wherever
  - Repeat: move to the best neighboring state.
  - If no neighbors better than current, quit.
  - Neighbors = small perturbations of w.
- Property
  - Many local optima

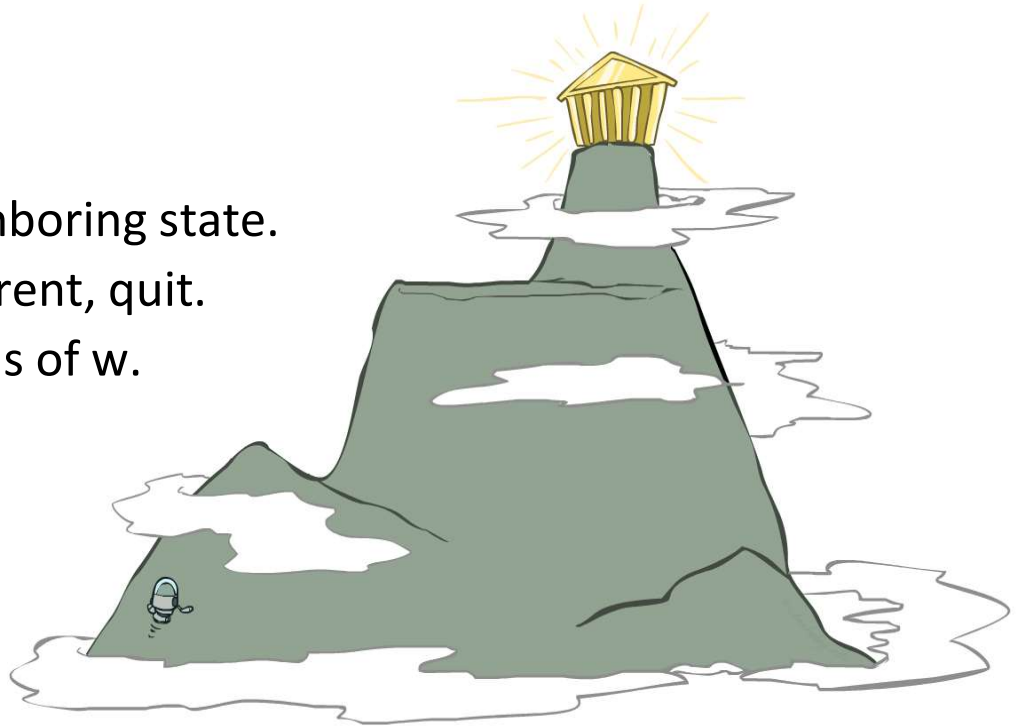

--> How to find a good local optimum?

Auto-Encoder (Crude Idea Sketch)

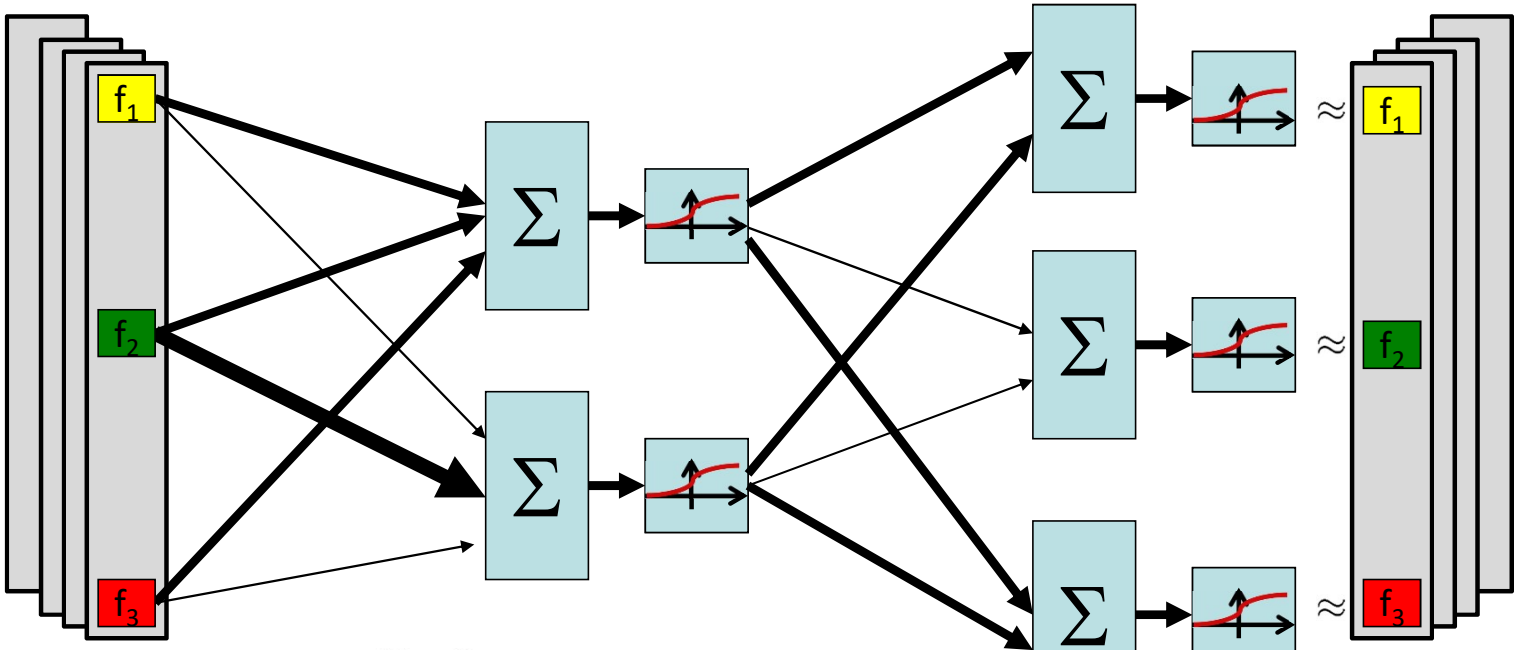

$$\min_{w} \sum_{i=1}^{m} \sum_{j=1}^{n} \left(f_j^{(i)} - h_j(w, f^{(i)}) \right)^2$$

Training Procedure: Stacked Auto-Encoder

- Auto-encoder
  - Layer 1 = “compressed” version of input layer
- Stacked Auto-encoder
  - For every image, make a compressed image (= layer 1 response to image)
  - Learn Layer 2 by using compressed images as input, and as output to be predicted
  - Repeat similarly for Layer 3, 4, etc.
- Some details left out
  - Typically in between layers responses get agglomerated from several neurons (“pooling” / “complex cells”)

Final Result: Trained Neural Network

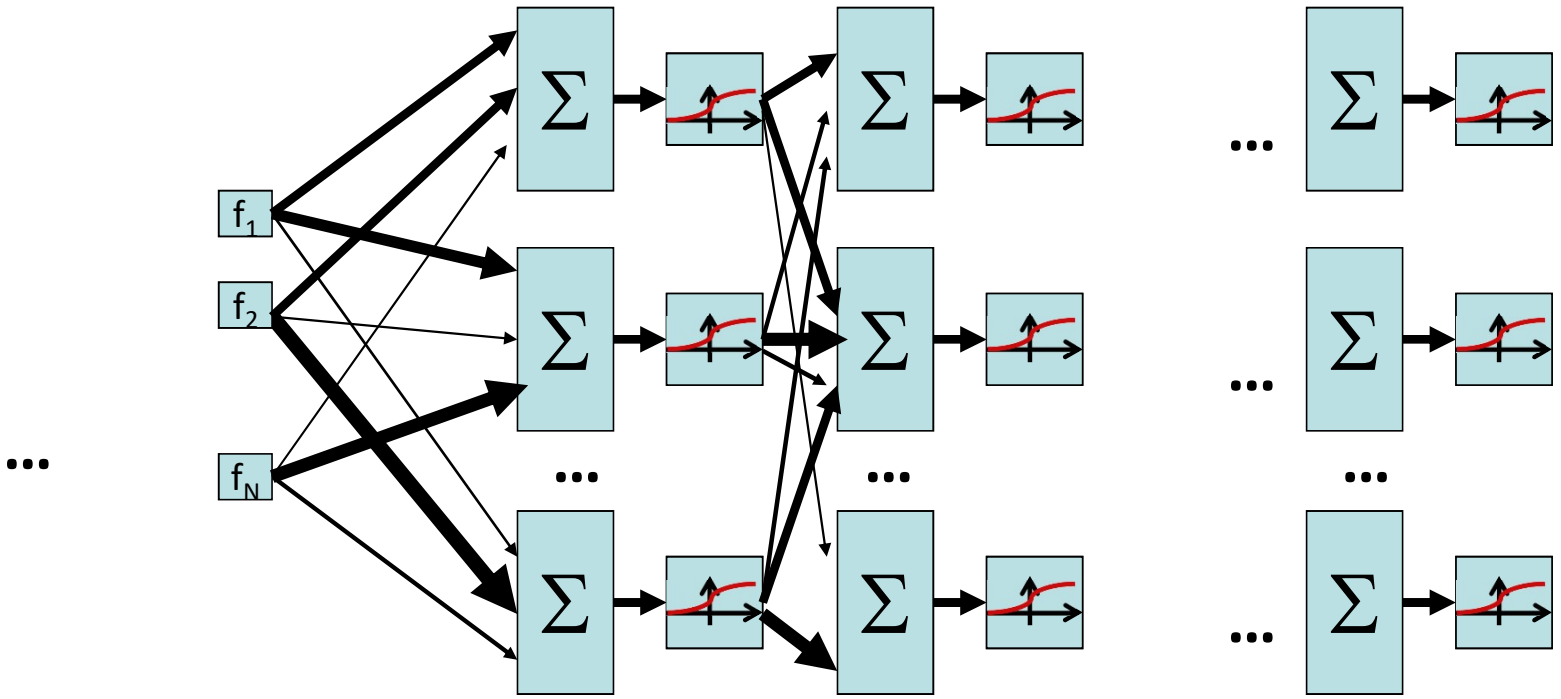

Final Result: Trained Neural Network