Artificial Intelligence and the Vector Institute

Geoffrey Hinton
Google Brain Team Toronto
& The Vector Institute

The sixty year battle for the soul of AI

The logic-inspired approach:

The essence of intelligence is using symbolic rules to manipulate symbolic expressions.

The biologically-inspired approach:

The essence of intelligence is learning the strengths of the connections in a neural network.

Highlights of the battle (1955-2015)

- 1960s: Very simple neural nets that only have one layer of features are wiped out by symbolic AI.
- 1980s: The back-propagation procedure allows neural networks to design their own features and to have multiple layers of features.
  - Lots of hype plus a few real applications.
- 1990s: On modest-sized datasets other machine learning methods work better. Oh dear! Neural networks wiped out again.
- 2009-2015: Deep neural nets work amazingly well if given a lot of data, a lot of compute power and a few technical tricks.
  - Google, Facebook, Microsoft, Apple all bet the farm.
  - Huge demand for well-trained students.

How to make a computer do what you want

- **The old way:** Write a program that tells the computer exactly what to do.

- **The new way:** First tell the computer to pretend to be a neural network.
  - Then show the neural network what to do by giving it many examples of inputs and outputs.

A close-up of a child holding a stuffed animal.

Input is an image

Output is a caption

A simplified model of a neuron

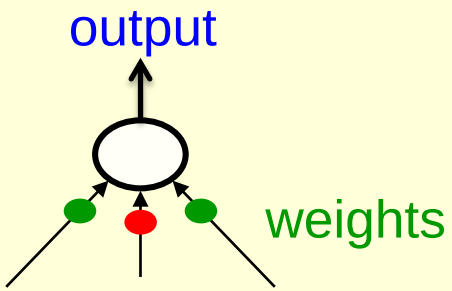

inputs coming from other neurons or sensors

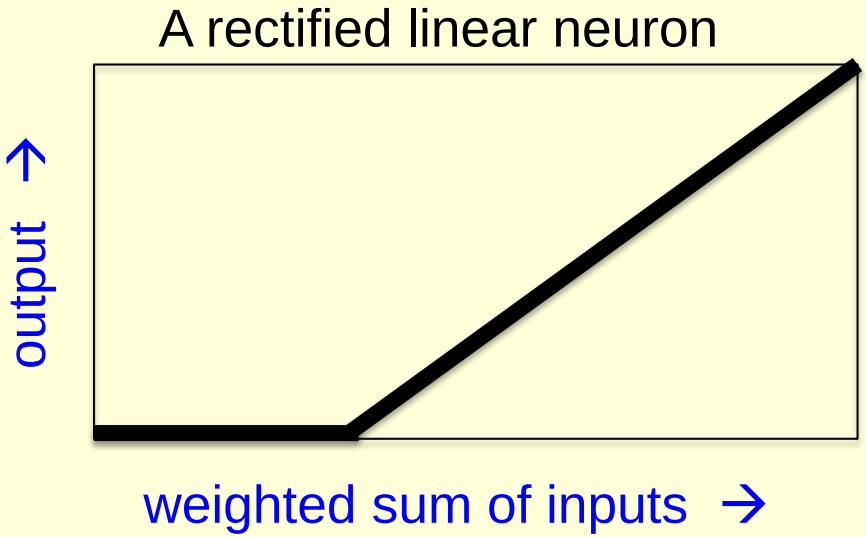

What is an artificial neural network?

- Arrange the neurons in layers

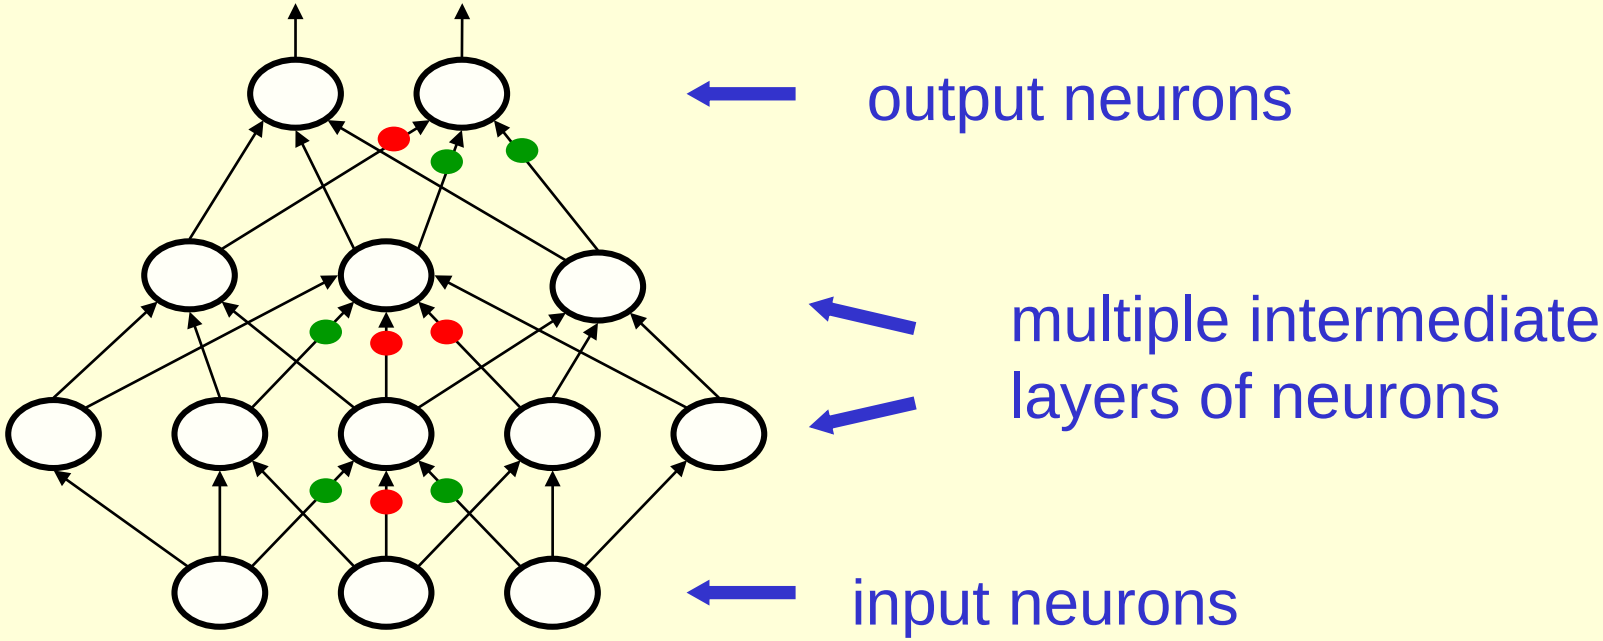

An inefficient way to train a neural network

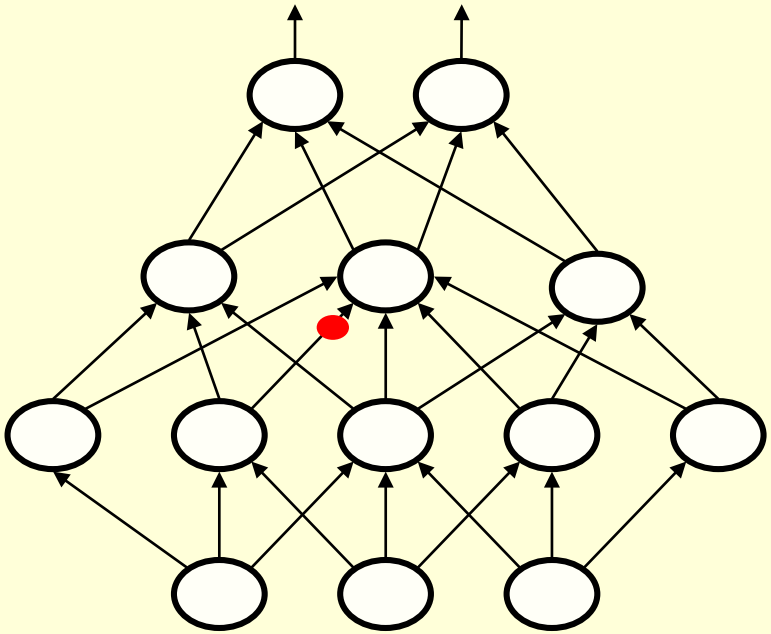

- Measure how well the network does on a set of examples.
- Pick one of the weights.
- Change the weight slightly and measure how well the network does.
- If the change helped, keep it.

An efficient way to train a neural network

Back-propagate the error signal to **compute** the effects of changing weights.

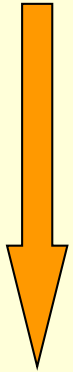

Compare the outputs with the correct answer to get error signal

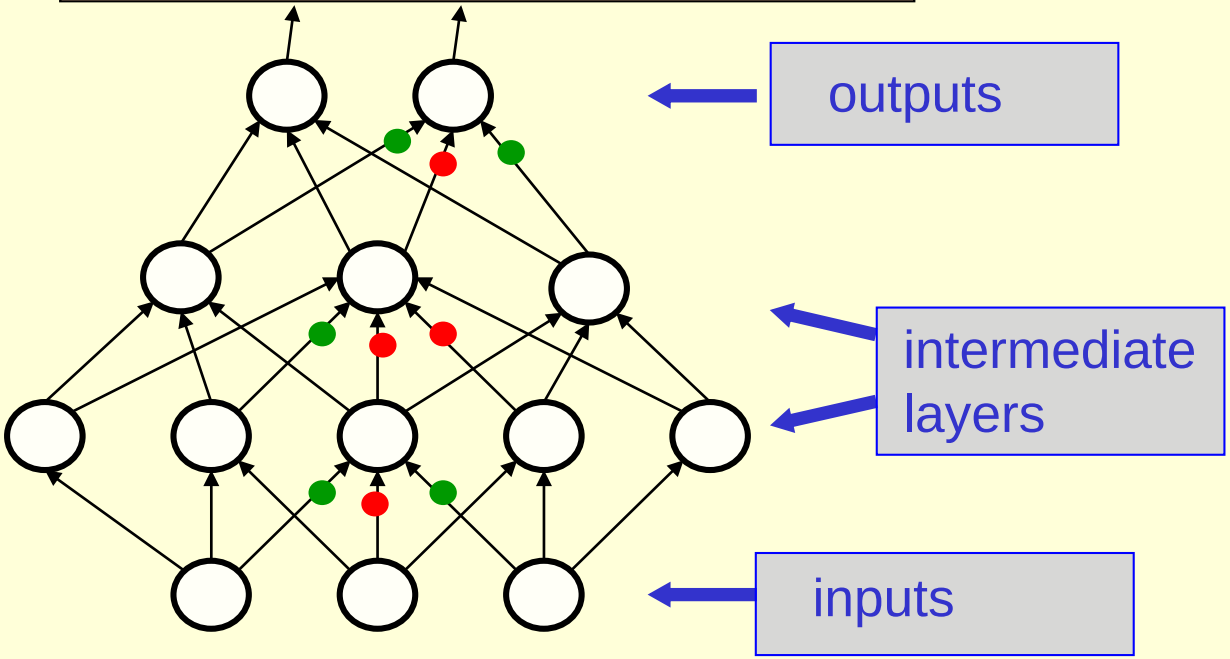

How research in Canada precipitated the AI revolution

- 1987 – 2017: Government funded basic research (NSERC & CIFAR)
- 2009: Speech recognition breakthrough (U. Toronto)
- 2012: Image classification breakthrough (U. Toronto)
- 2014: Machine translation breakthrough (U. Montreal)

Machine Translation

- This is a perfect problem for symbolic AI because the input and output are strings of symbols.
  - Also, linguists already know lots of rules.

- An utterly crazy neural network approach:
  - Learn to convert the input sentence into a big pattern of neural activities that is language independent.
  - This pattern is a thought vector.
  - Learn to convert the thought vector into a sentence in the target language.

What is a thought?

- It is not anything like a symbolic expression.

- It is a big neural activity pattern.
  - We refer to it by using a symbol string that causes it (or that it causes).
  - But the way we refer to it is quite different from what it is.

- Sentences are not even the shadows of thoughts.

Some things that will be coming soon

- Most medical images will be read by neural networks.
  - An app on your phone will diagnose skin cancer
- We will get personal assistants that are really smart.
  - and they will know how to hold a conversation.
- Any organization with a lot of data will get much better at predicting things it cares about.
  - It will use neural networks running in the cloud.

The Vector Institute

- Tomi Poutanen, Jordan Jacobs and Prof. Rich Zemel decided to raise money for an independent machine learning institute.
  - They recruited Ed Clarke, and other U of T machine learning professors to help.

- With support from on high at U of T they quickly raised:

  50M from Ontario

  ~40M from the federal government (share of a pan-Canadian initiative)

  ~80M from Industry (with almost no strings attached!)

  ~20M more from Ontario for training

The goals of the Vector Institute

- Keep Toronto at the forefront of basic research in the new AI (stop the brain drain).
- Develop game-changing applications of the new AI (with special emphasis on healthcare).
- Educate the many highly qualified people that Ontario industry needs. (more profs and professional masters)
- Encourage start-ups and scale-ups and provide expertise for local companies.

How will the money be spent?

- Employ world-class research scientists who have courtesy appointments at U of T (or other universities).
- Pay all or part of the salaries for more professors in various departments that would like to have them.
- Provide extra stipends to attract and retain top talent.
  (NYT says new PhD's in deep learning get 300-500K in industry).
- Provide support for many more grad students.
- Rent space at MARS.

THE END