Neuroscience-informed Artificial Intelligence

Andrew Ng

Artificial Intelligence

The brain's algorithm

Some suggestion that the mammalian brain (neocortex) may use essentially *the same algorithm* to understand many different input modalities. (e.g., Fukushima, Hinton, Hawkins, etc.)

- Example: Ferret experiments, in which the "input" for vision was plugged into auditory part of brain, and the auditory cortex learns to "see." [Roe et al., 1992]
- Example: Sensory remapping in humans.
  - Visual cortex used by blind persons for touch.
  - Tapping out images on tongue, which is then used to "see."

Can we discover or approximate the brain's learning algorithm, and build a small piece of an "artificial human brain"?

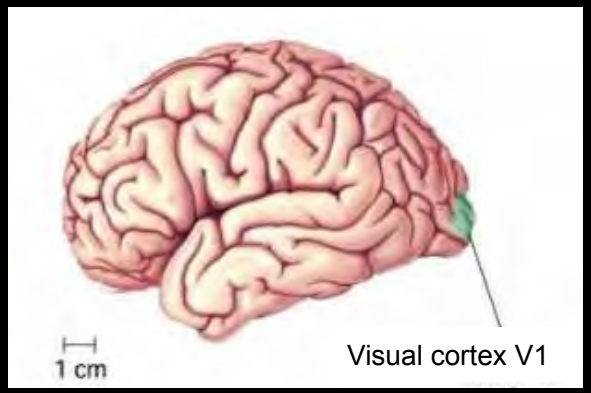

Artificial Intelligence

Is it feasible?

Current work

Example: Visual cortical area V1

To what extent can learning algorithms today mimic computations in the brain?

V1 is the first stage of the visual cortex.

Known to act as “edge detectors.”

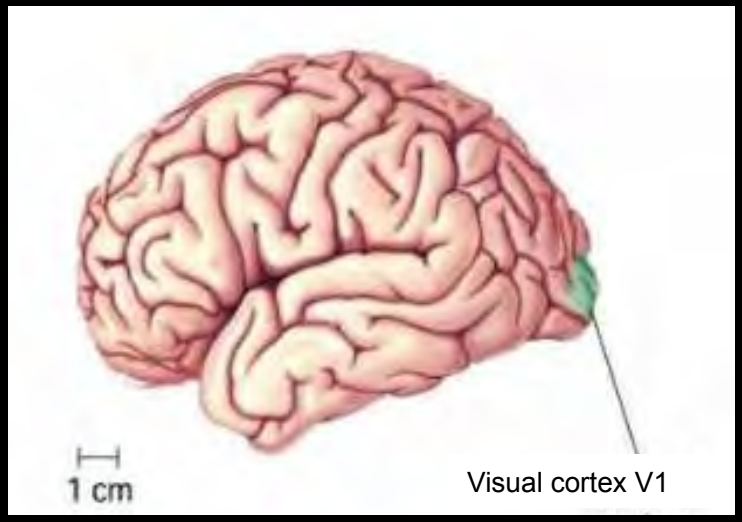

Visual cortex V1

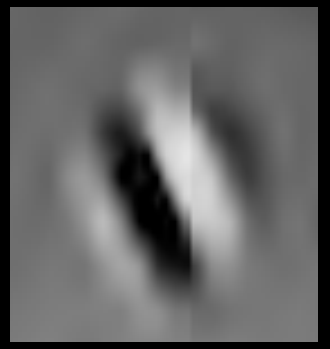

V1: “Edge detectors.”

[DeAngelis, Ohzawa & Freeman, 1995]

Model of V1

Learning algorithm for V1.

Interpret as finding a “sparse code” of the input image. (Olshausen & Field,1996)

Example:

$x \approx 0.8 \cdot \phi_{87} + 0.3 \cdot \phi_{376} + 0.5 \cdot \phi_{411}$

Decompose an input image x into a sum of simpler “basic images” ϕ_i .

Example of learned bases

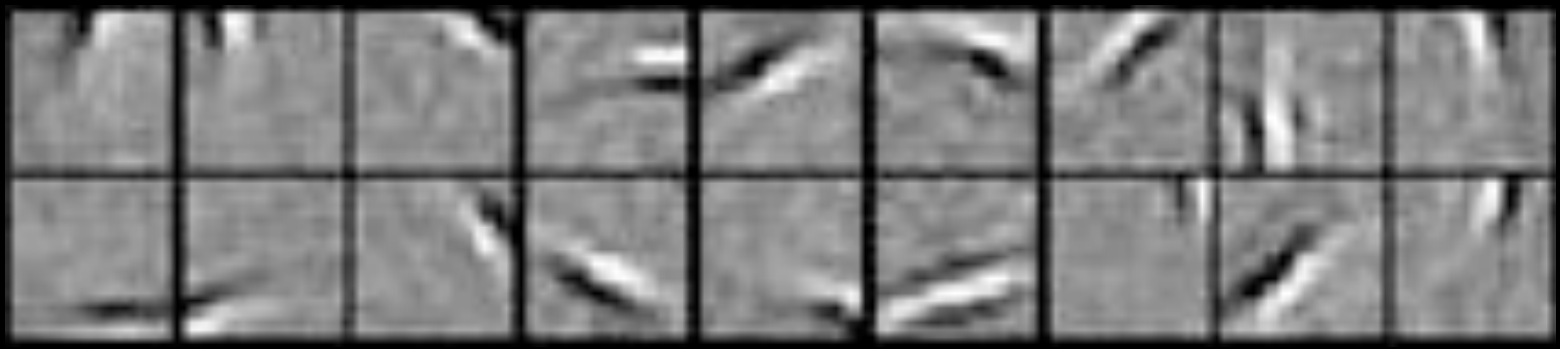

Examples of basis patches $\phi_i \in \mathbb{R}^{n \times n}$ learned. Many basis looks like edge detectors.

• Method hypothesizes that edge-like patches are the most “basic” elements of a scene, and represents an image in terms of the edges that appear in it.

• Algorithm has “invented” edge detection.

• Learned model corresponds fairly closely on many dimensions to measurements of V1. (van Hataren & van de Schaaf, 1998)

Learning from audio

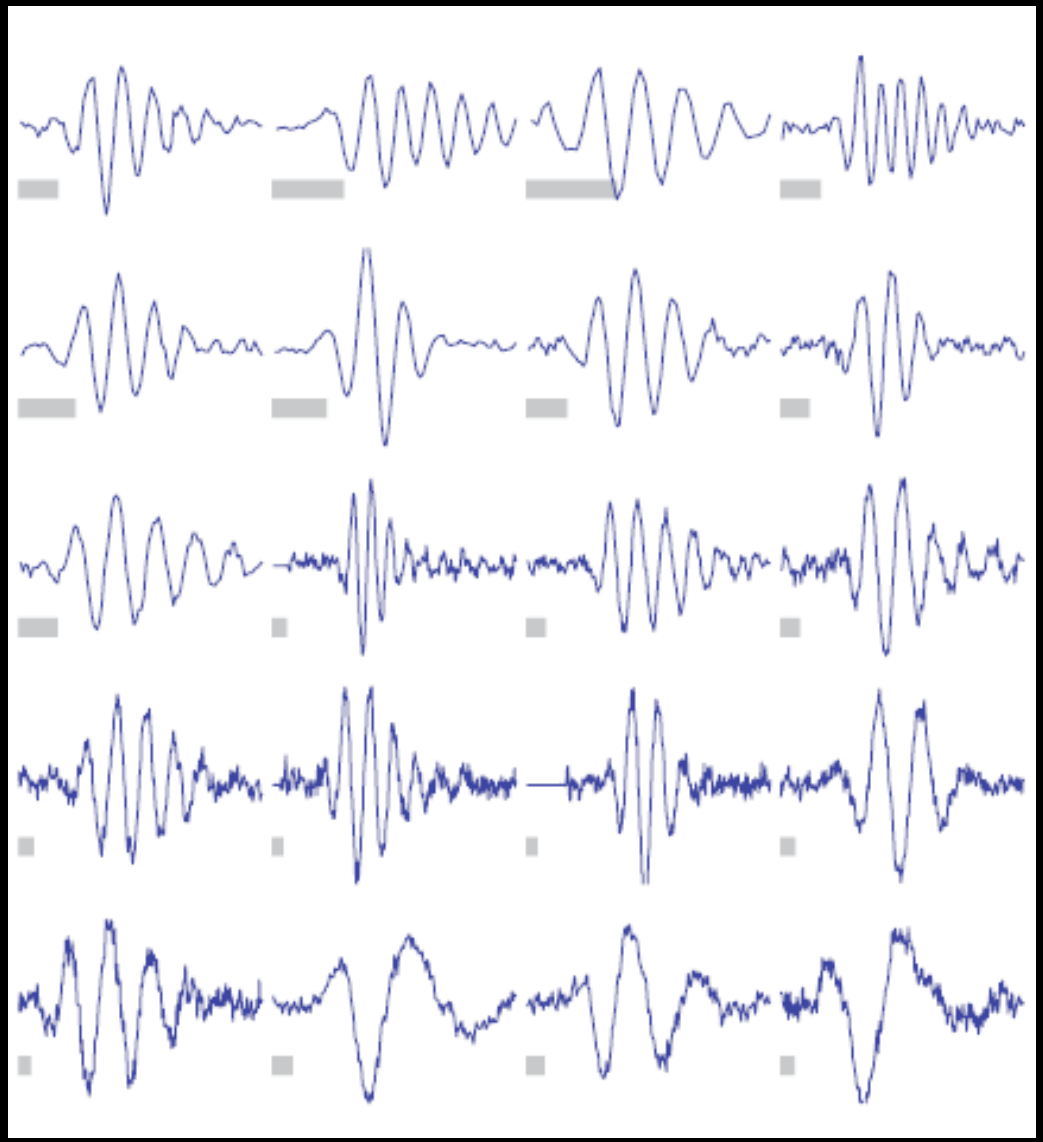

[Evan Smith & Mike Lewicki, 2006]

Learning from audio

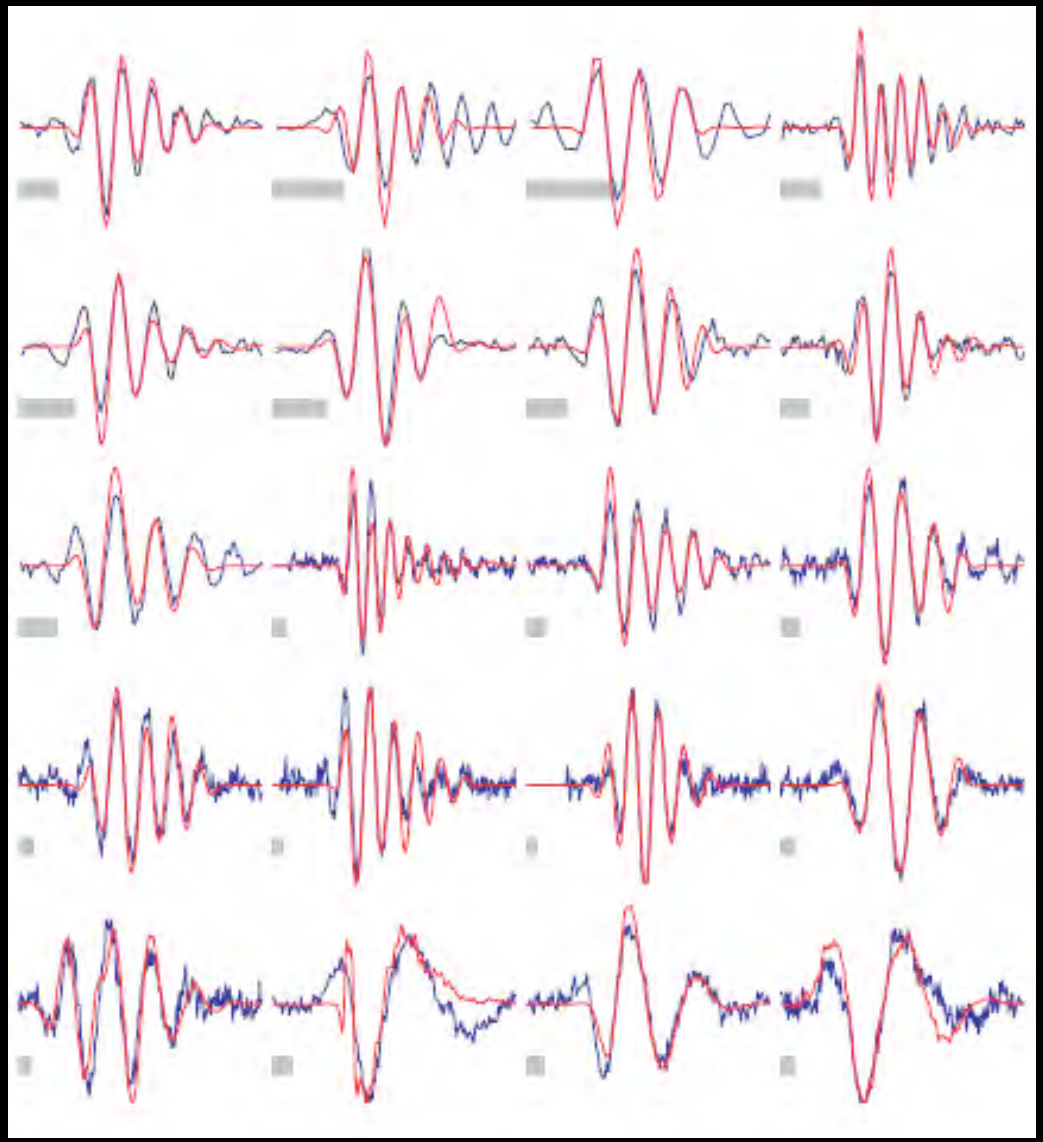

[Evan Smith & Mike Lewicki, 2006]

... and much more.

Community

Academic:
- Growing interest in machine learning/AI communities.
- Existing academic workshops/meetings, etc.

Government.

Industry: IBM cognitive computing initiative. Numenta.

Popular imagination: Singularity (Kurzweil), On Intelligence (Hawkins & Blakeslee).

Didn’t we do this before?

Brain's learning algorithm

An order-of-magnitude argument:

Your brain has 10^{14} synapses (connections).

You'll live for 10^{9} seconds.

If each synapse requires just 1 bit to parameterize, you need to learn 10^{14} bits in 10^{9} seconds.

That's 10^{5} bits/second.

Most of human learning is *unsupervised*.

(Geoff Hinton, pers. comm.)

Basic science goal: Understanding the brain

Engineering goal: Neuroscience-informed Artificial

The AI dream