COMP9444
Neural Networks and Deep Learning

1. Neuroanatomy

What is a Neural Network?

- massively parallel distributed processor made up of simple processing units
- knowledge acquired from environment through a learning process
- knowledge stored in the form of synaptic weights

Why Neural Networks?

- biologically inspired
- good learning properties
- continuous, nonlinear
- well adapted to certain tasks
- fault tolerant
- graceful degradation

Sub-Symbolic Processing

Theories about Intelligence

- 380BC Plato (Rationalism - innateness)
- 330BC Aristotle (Empricism - experience)
- 1641 Descartes (mind-body Dualism)
- 1781 Kant (Critique of Pure Reason)
- 1899 Sigmund Freud (Psychology)
- 1953 B.F. Skinner (Behaviourism)

Artificial Intelligence Origins

- 1642 Blaise Pascal (mechanical adding machine)
- 1694 Gottfried Leibniz (mechanical calculator)
- 1769 Wolfgang von Kempelen (Mechanical Turk)
- 1837 Charles Babbage & Ada Lovelace (Difference Engine)
- 1848 George Boole (the Calculus of Logic)
- 1879 Gottlob Frege (Predicate Logic)
- 1950 Turing Test
- 1956 Dartmouth conference

Neural Network Origins

- 1943 McCulloch & Pitts (neuron models)
- 1948 Norbert Wiener (Cybernetics)
- 1948 Alan Turing (B-Type Networks)
- 1955 Oliver Selfridge (Pattern Recognition)
- 1962 Hubel and Wiesel (visual cortex)
- 1962 Frank Rosenblatt (Perceptron)

Serial Symbolic AI

- 1956 Newell & Simon (Logic Theorist)
- 1959 John McCarthy (Lisp)
- 1959 Arther Samuel (Checkers)
- 1965 Joseph Weizenbaum (ELIZA)
- 1967 Edward Feigenbaum (Dendral)

Neural Network "Dark Ages"

- 1969 Minsky & Papert published Perceptrons, emphasizing the limitations of neural models, and lobbied agencies to cease funding neural network research.
- from 1969 to 1985 there was very little work in neural networks or machine learning.
- a few exceptions, e.g. Stephen Grossberg, Teuvo Kohonen (SOM), Paul Werbos.

Knowledge-Based Systems

- 1970s and early 1980s, AI research focused on symbolic processing, Expert Systems
- Some commercial success, but ran into difficulties:
  - combinatorial explosion in search spaces
  - difficulty of formalising everyday knowledge as well as expert knowledge

Neural Network Renaissance

- 1986 Rumelhart, Hinton & Williams (multi-layer, backprop)
- 1989 Dean Pomerleau (ALVINN)
- late 1980's renewed enthusiasm, hype
- 1990s more principled approaches
- 2000's SVM, Bayesian models became more popular
- 2010's deep learning networks, GPU's
- 2020's spiking networks(?)

Applications of Deep Learning

- Image processing
  - classification
  - segmentation
- Language processing
  - translation
  - semantic disambiguation
  - sentiment analysis
- Combining images and tex
  - automatic captioning
- Game playing
  - AlphaGo
  - Deep Q-Learning

History of Deep Learning

Two perspectives on the history of Deep Learning

Viewpoint 1: Focusing on recent work (after 2012)

`https://www.cs.toronto.edu/~hinton/absps/NatureDeepReview.pdf`

Viewpoint 2: Focusing on earlier work (before 2012)

`http://people.idsia.ch/~juergen/deep-learning-overview.html`

Neuroanatomy

- Central Nervous System
  - Brain
  - Spinal cord
- Peripheral Nervous System
  - Somatic nervous system
  - Autonomic nervous system
  - Enteric nervous system

Brain Regions

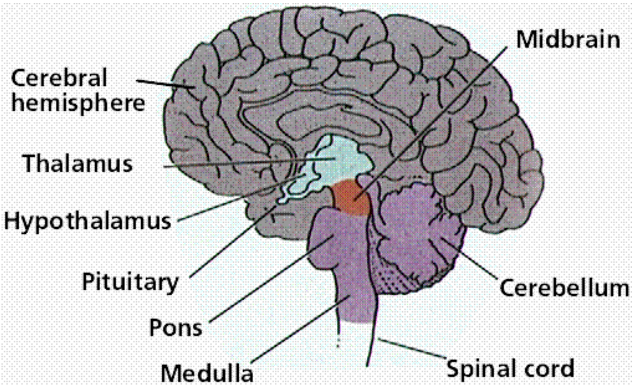

Cerebral Cortex

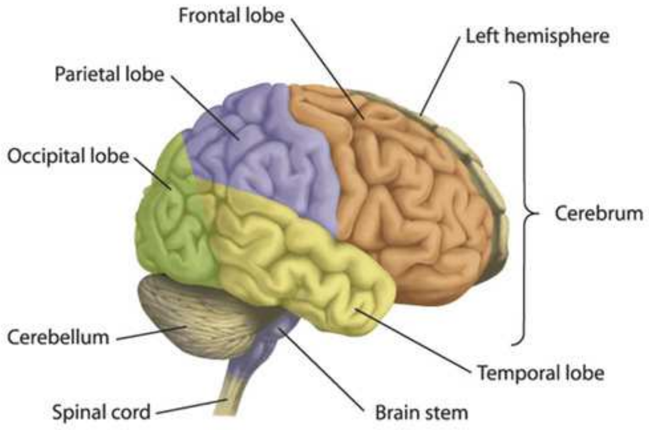

Cerebral Cortex

- "cortex" from Latin word for "bark" (of tree)
- cortex is a sheet of tissue making up outer layers of brain, 2-6cm thick
- right and left sides connected by corpus callosum
- functions: thought, voluntary movement, language, reasoning, perception

Brain Stem

- general term for area of brain between the thalamus and spinal cord
- includes medulla, pons, tectum, reticular formation and tegmentum
- functions: breathing, heart rate, blood pressure, and others

Cerebellum

- from Latin word for "little brain"
- functions: movement, balance, posture

Midbrain

- functions: vision, audition, eye movement, body movement

Thalamus

- receives sensory information and relays it to the cerebral cortex
- also relays information from the cerebral cortex to other areas of the brain, and the spinal cord
- functions: sensory integration, motor integration

Hypothalamus

- composed of several different areas at the base of the brain
- the size of a pea (about 1/300 of the total brain weight)
- functions: body temperature, emotions, hunger, thirst, circadian rhythms

Limbic System

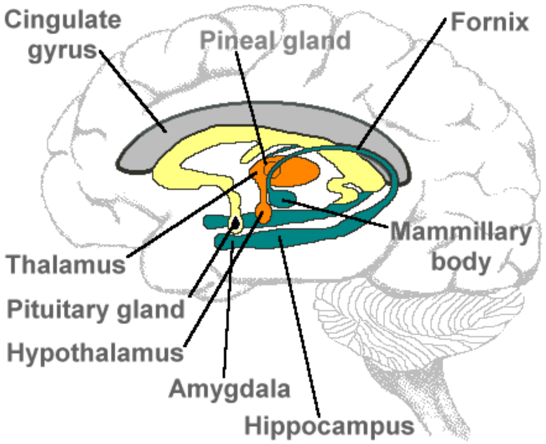

Limbic System

- group of structures including amygdala, hippocampus, mammillary bodies and cingulate gyrus
- important for controlling the emotional response to a given situation
- hippocampus also important for memory
- functions: emotional behaviour

Brain Functions

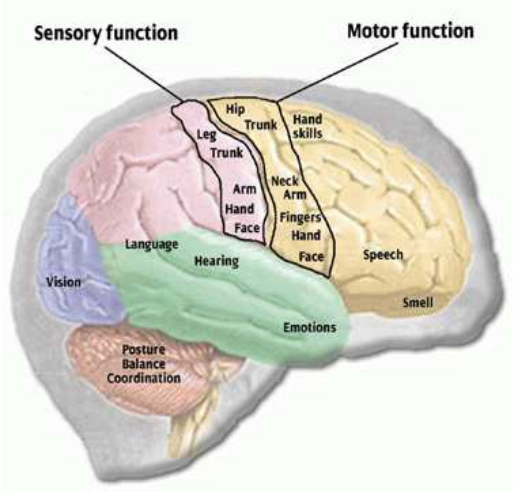

Neurons as Body Cells

- The body is made up of billions of cells. Cells of the nervous system, called neurons, are specialized to carry “messages” through an electrochemical process.
- The human brain has about 100 billion neurons, and a similar number of support cells called “glia”.
- Neurons are similar to other cells in the body in some ways, such as:
  - neurons are surrounded by a cell membrane
  - neurons have a nucleus that contains genes (DNA)
  - neurons carry out basic cellular processes like protein synthesis and energy production

Structure of a Typical Neuron

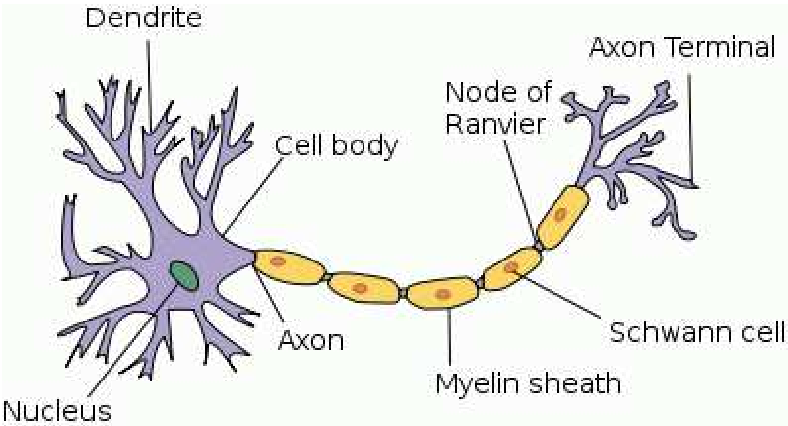

Neurons versus Body Cells

- Neurons have specialized extensions called dendrites and axons Dendrites bring information to the cell body, while axons take information away from the cell body.
- The axon of one neuron can connect to the dendrite of another neuron through an electrochemical junction called a synapse.
- Most neurons have only one axon, but the number of dendrites can vary widely:
  - Unipolar and Bipolar neurons have only one dendrite
  - Purkinje neurons can have up to 100,000 dendrites

Variety of Neuron Types

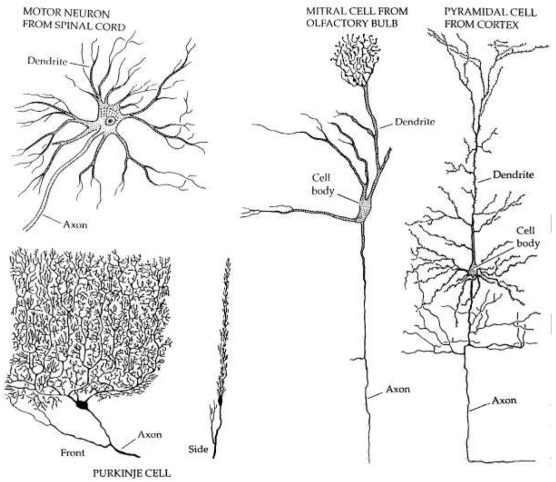

Axons and Dendrites

- Dendrites are typically less than a millimetre in length
- Axons can vary in length from less than a millimetre to more than a metre (motor neurons)
- Long axons are sometimes surrounded by a myelinated sheath, which prevents the electrical signal from dispersing, and allows it to travel faster (up to 100 m/s).

Synapse

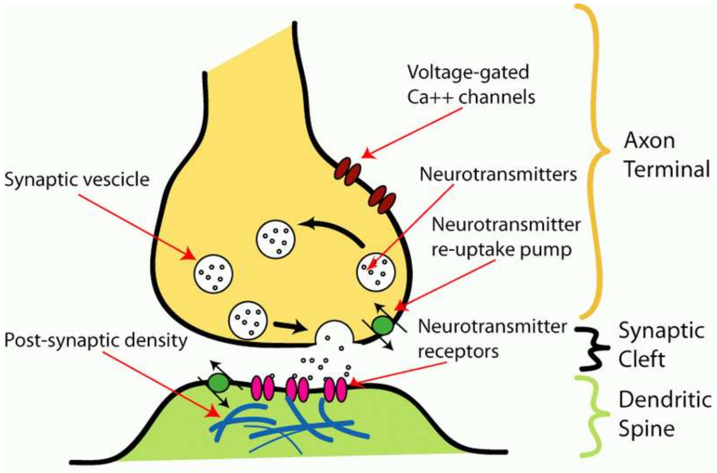

Synapses and Ion Channels

- electrical pulse reaches the endbulb and causes the release of neurotransmitter molecules from little packets (vesicles) through the synaptic membrane
- transmitter then diffuses through the synaptic cleft to the other side
- when the neurotransmitter reaches the post-synaptic membrane, it causes a change in polarisation of the membrane
- the change in potential can be excitatory (moving the potential towards the threshold) or inhibitory (moving it away from the threshold)

Ion Channel

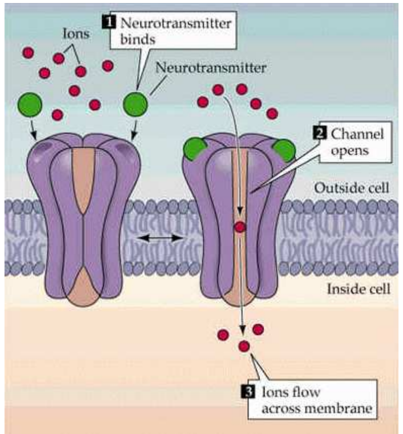

The Big Picture

- human brain has 100 billion neurons with an average of $10,000$ synapses each
- latency is about 3-6 milliseconds
- therefore, at most a few hundred “steps” in any mental computation, but massively parallel

Hubel and Weisel – Visual Cortex

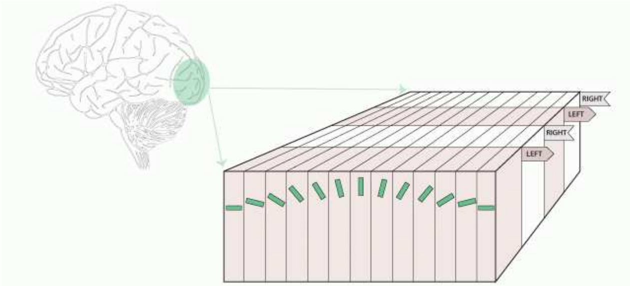

- cells in the visual cortex respond to lines at different angles
- cells in V2 respond to more sophisticated visual features
- Convolutional Neural Networks are inspired by this neuroanatomy
- CNN’s can now be simulated with massive parallelism, using GPU’s

Convolutional Networks

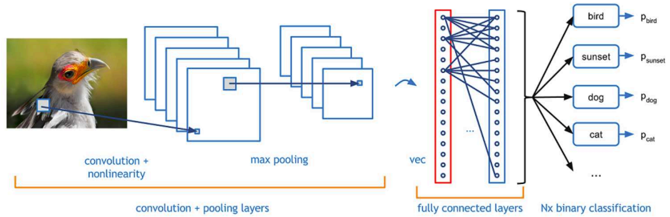

Suppose we want to classify an image as a bird, sunset, dog, cat, etc.

If we can identify features such as feather, eye, or beak which provide useful information in one part of the image, then those features are likely to also be relevant in another part of the image.

We can exploit this regularity by using a convolution layer which applies the same weights to different parts of the image.

Convolutional Filters

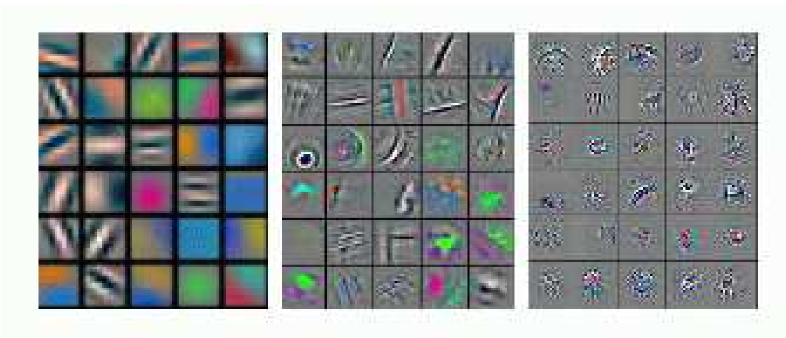

First Layer Second Layer Third Layer

Recurrent Neural Networks

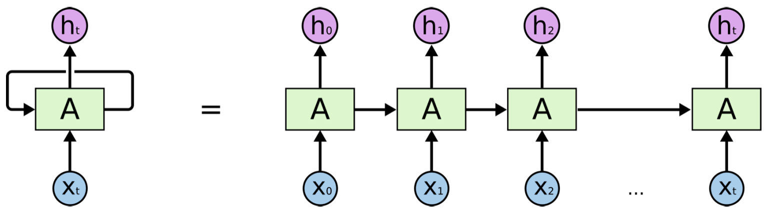

- can “unroll” a recurrent architecture into an equivalent feedforward architecture, with shared weights
- useful for processing language or other temporal sequences

Deep Q-Learning

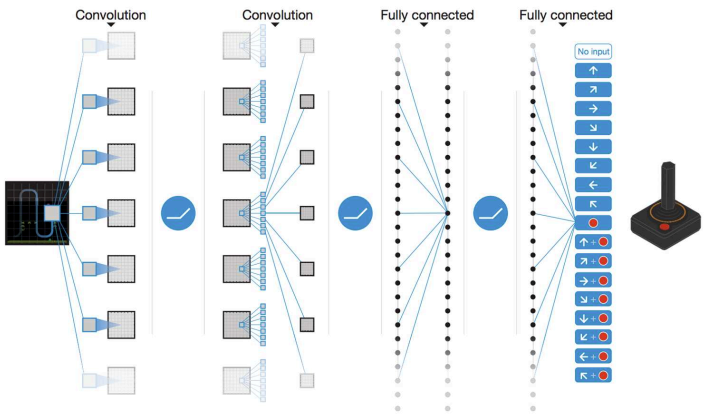

Autoencoder Networks

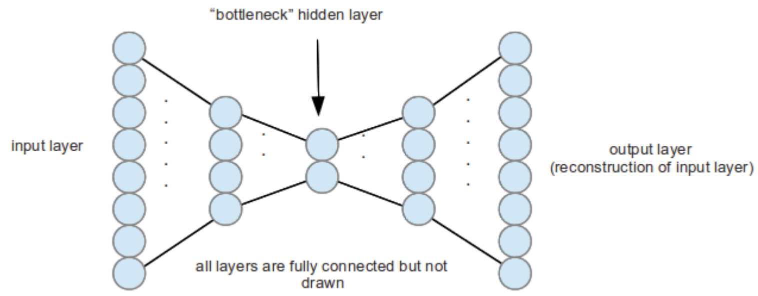

- output is trained to reproduce the input as closely as possible
- activations normally pass through a bottleneck, so the network is forced to compress the data in some way

Generative Adversarial Networks

Spiking Neurons

- biological neurons spike in different patterns (quiescent, persistent, sporadic)
- spike timing might carry important information
- most NN models ignore timing information, but some work has been done on spiking network models
- in the future, special hardware might lead to a revolution for spiking networks, similar to what GPU's provided for CNN's