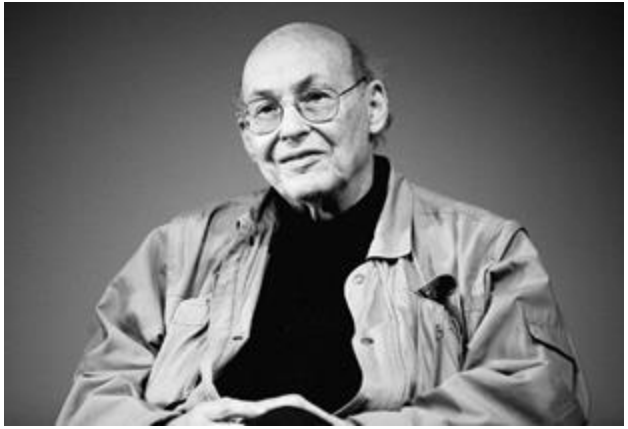
Introduction to Cognitive Science

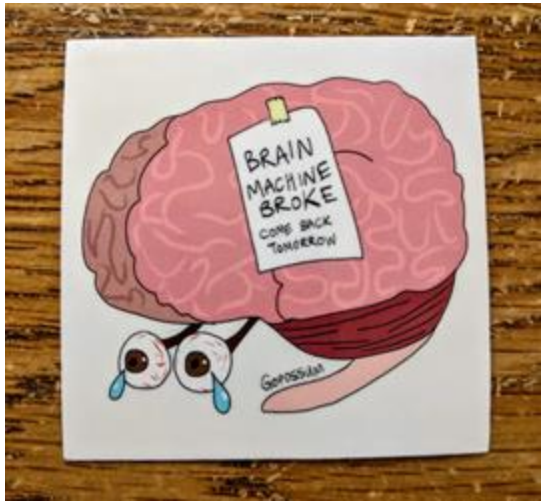


Dr. Mark Miller

MONASH University
UNIVERSITY OF TORONTO
CHAIN



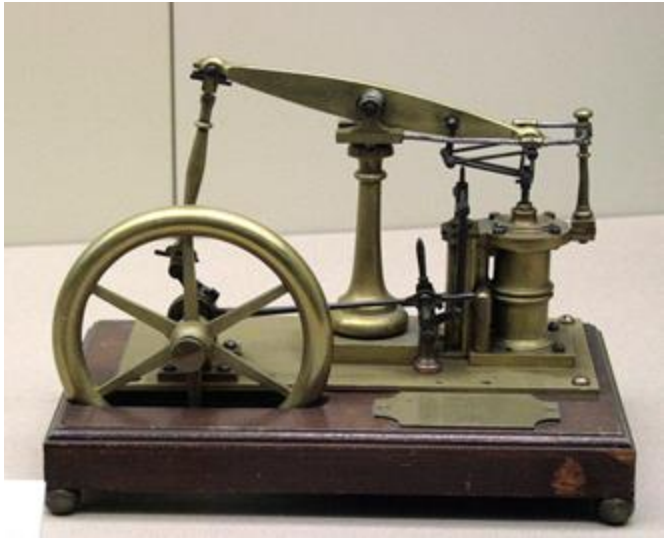




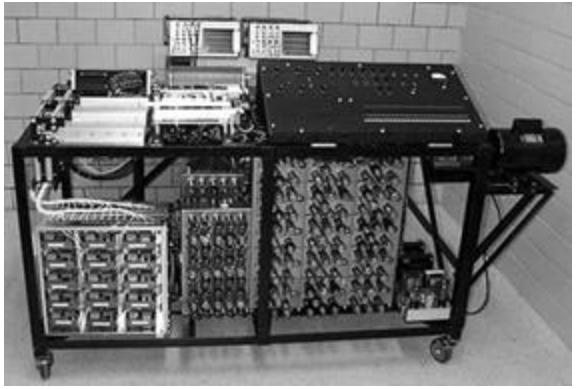
Minsky's description emphasizes that what matters isn't only (or primarily) the material (the "meat"), but how it's **organized** to support thinking - the machine part - the information processing itself.

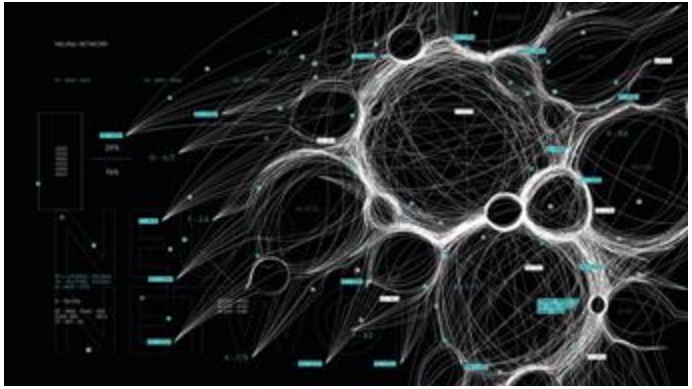


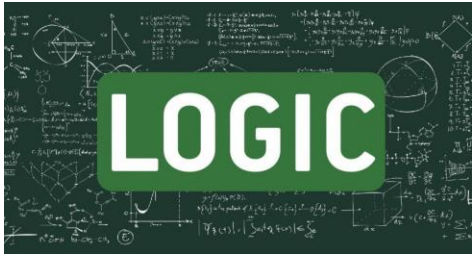




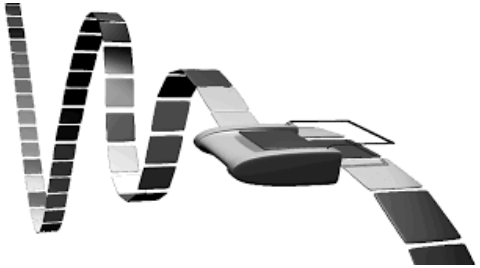


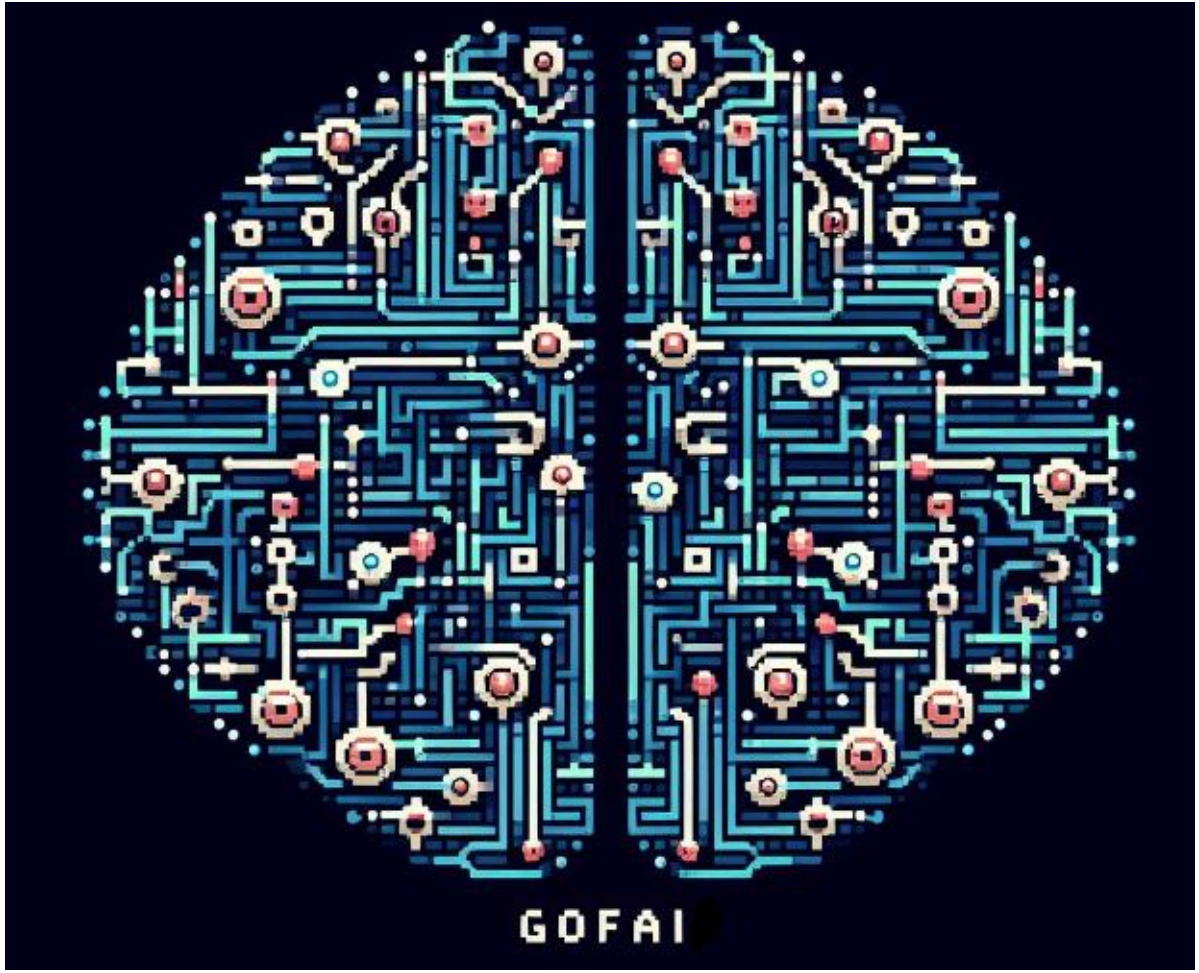


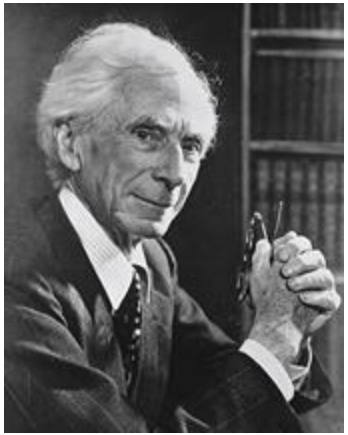



LOGIC





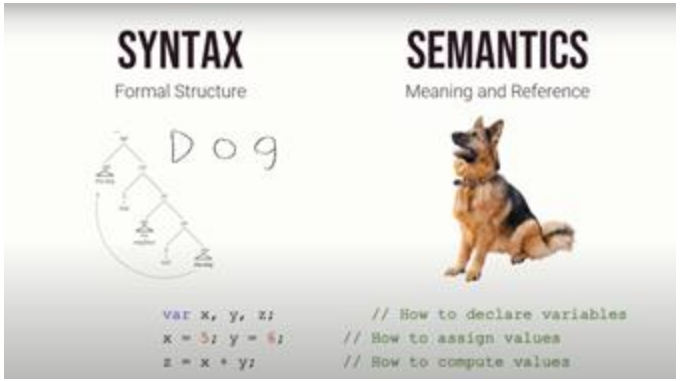

GOFAI

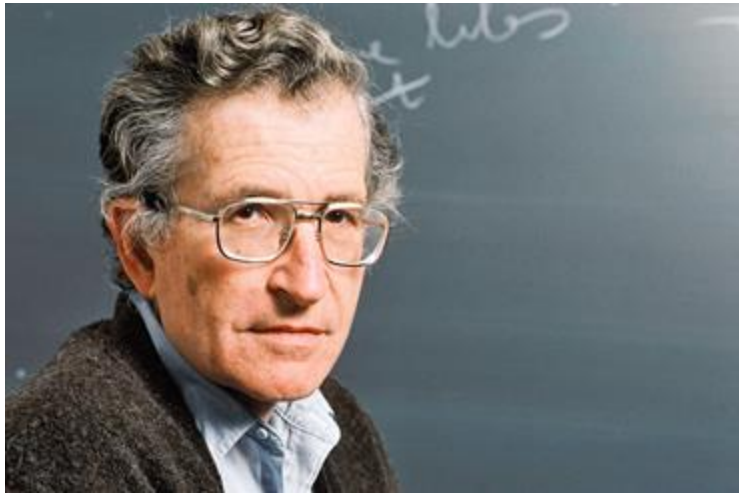






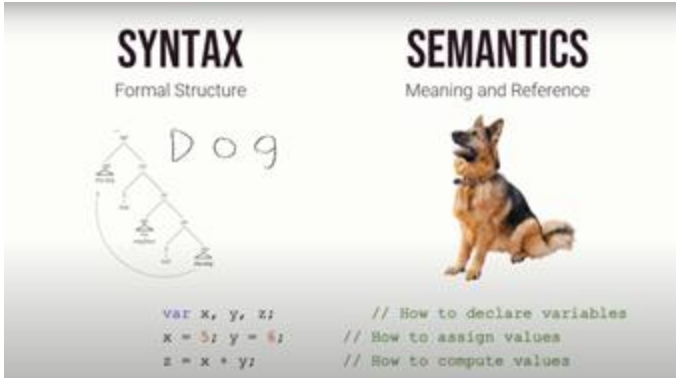
Formal logic changed how we approached symbols, moving away from their meaning (**semantics**) and focusing more on their form and structure (**syntax**). In this context, logic became about **manipulating abstract symbols according to rules**, regardless of what the symbols actually represent.







"A sentence was a string of symbols. A language was a set of such strings. A grammar was a set of recursive procedures for generating such sets. **Language was syntacticized**—placed mathematically within a Post system, with abstract symbols manipulated in algorithmic fashion by precise formal rules".





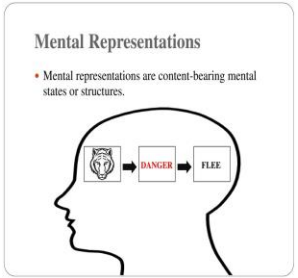
Functionalism is the idea that mental states are not defined by what they are made of, **but by what they do**

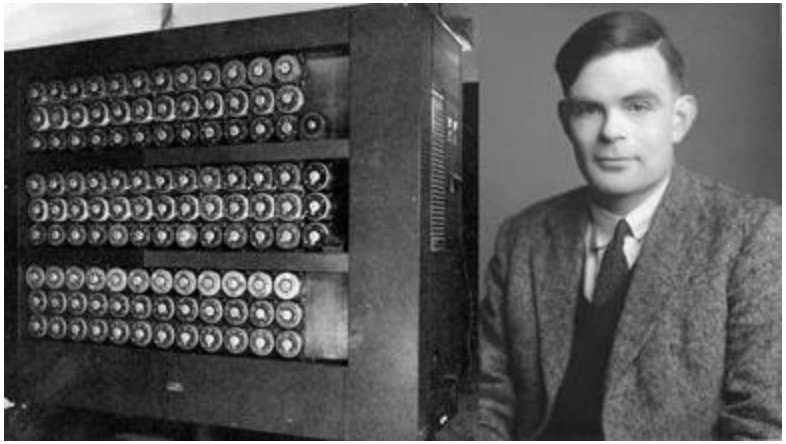


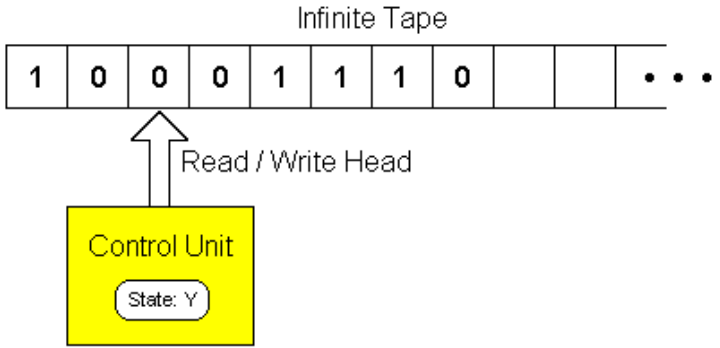


In this view, mental states are also **representational**. They represent the world and allow us to think about things that aren't immediately present

This ability to manipulate internal representations according to rules is precisely what we expect from both humans and computers. In this sense, the **Representational Theory of Mind** aligns with classical AI—human cognition and machine processes both operate by manipulating symbols according to rules.

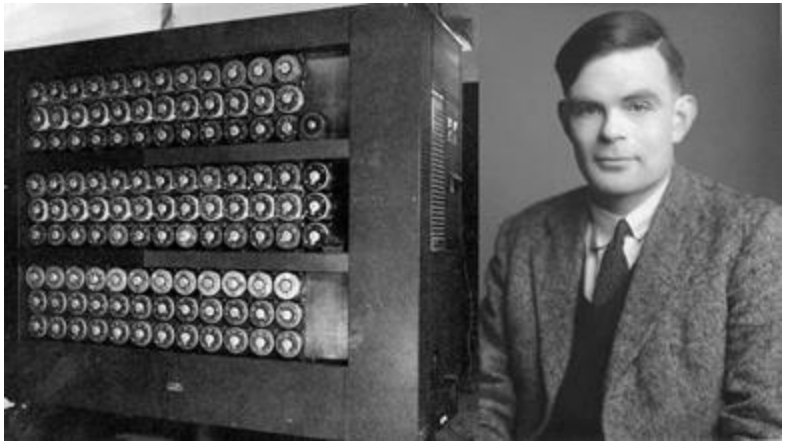


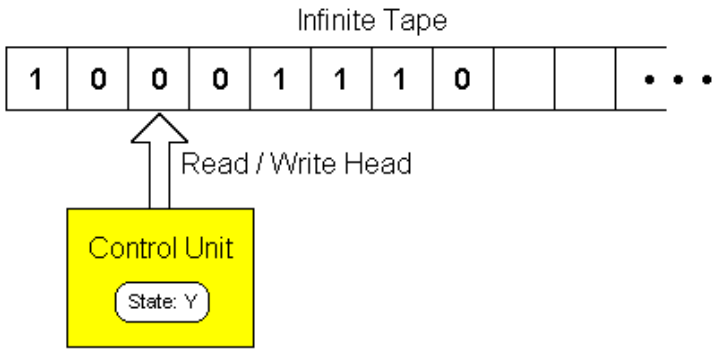




Alan Turing demonstrated in 1936 that a machine performing a sequence of simple computations could solve any well-defined problem.

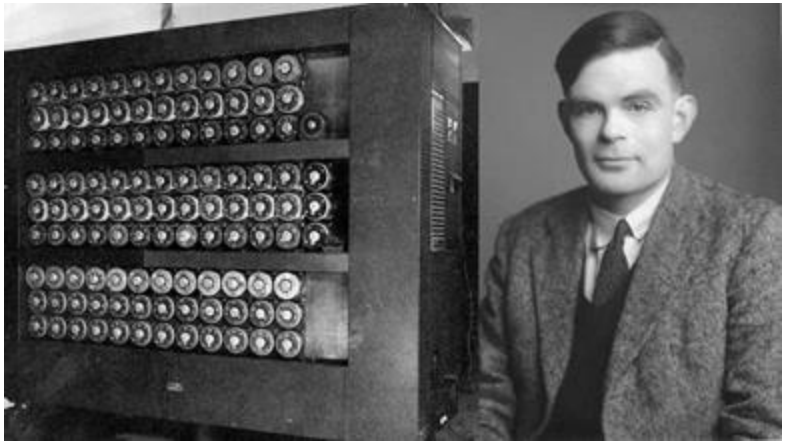
The machine reads symbols on the tape one at a time, follows predefined rules to either write a new symbol, move the tape left or right, or change its internal state. By following these simple rules, a Turing machine can simulate any algorithmic process, demonstrating how computers perform calculations.





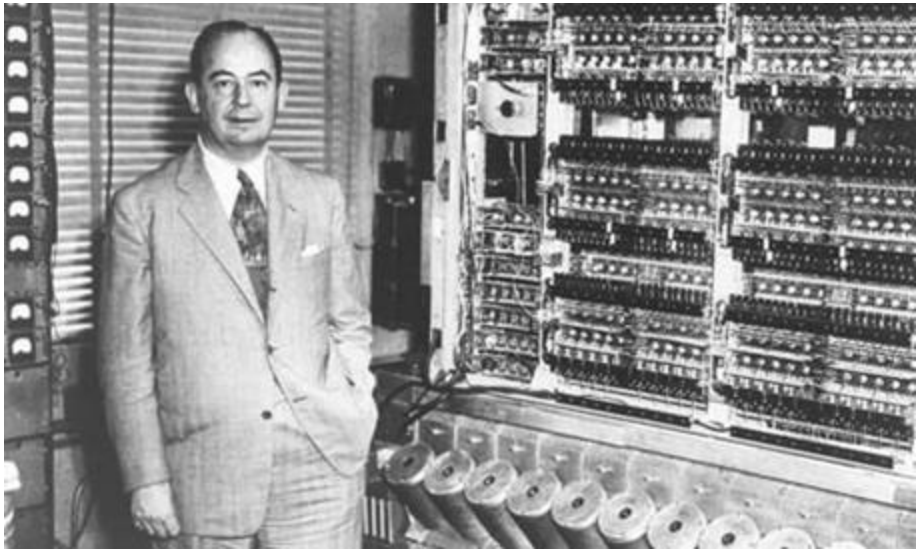
Alan Turing demonstrated in 1936 that a machine performing a sequence of simple computations could solve any well-defined problem.

Turing's thesis, which claims that if cognitive processes are **algorithmically calculable**, then they could be carried out by a Turing Machine



Alan Turing demonstrated in 1936 that a machine performing a sequence of simple computations could solve any well-defined problem.

The digital computer, designed by **John von Neumann** in the 1940s, became the model for this line of thinking. Just as software runs on hardware, cognition was thought to operate on the brain's physical substrate.




ANY
QUESTIONS?



GOFAI has had its share of **success stories**, most notably **Deep Blue**, the chess-playing AI that defeated world champion Garry Kasparov in 1997.

More recently, AI models like **ChatGPT** demonstrate how symbol manipulation on a massive scale can simulate human-like language





“In this view, "mindware" is like software—dependent on the physical brain but abstract enough to be understood in terms of **computation** rather than biology. The key idea is that cognition is a computational process, one that can be studied and understood through the lens of formal systems and algorithms, rather than the specific biological details of the brain”.



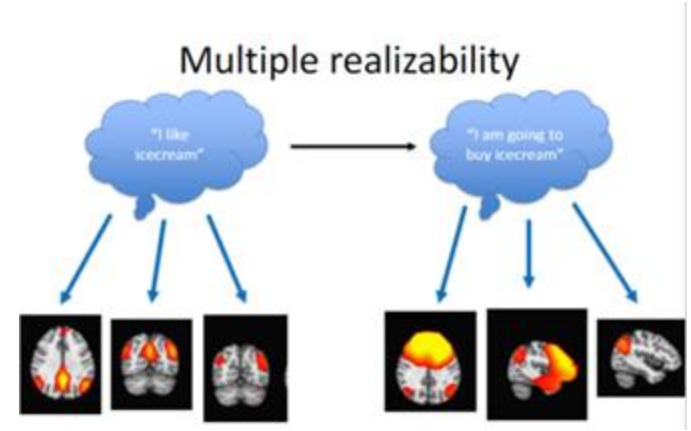


Churchland and Sejnowski (1992) suggest that something is considered a computer when its inputs and outputs can be understood as representing parts of a function that we care about. For example, a pocket calculator computes because its buttons and display can be interpreted as numbers, and its internal processes match meaningful changes in arithmetic





Dave Chalmers adds that a device performs computation if its physical state changes follow the steps of a formal set of rules. This idea can even apply to natural systems. For instance, if an animal's behavior changes based on how many predators it detects, nature has, in a way, created a basic computer..


Multiple realizability
"I like icecream"
"I am going to buy icecream"


© BNPS.CO.UK






ANY
QUESTIONS?



Now ask yourself; does John **understand Chinese**?



“If thoughts are composed out of symbols, then from **where to those symbols get their meanings**?...

So the problem is this: once we’ve stripped out meanings from the foundations of our mechanical account of cognition, ***how do they get back in again***?”



First, the thought experiment does not show that machines cannot think – Searle clarifies his view: “On the contrary, the brain is a machine and the brain can think” and in his 1980 paper, in response to the question ‘Could a machine think?’ he writes “The answer is, obviously, yes. We are precisely such machines.”

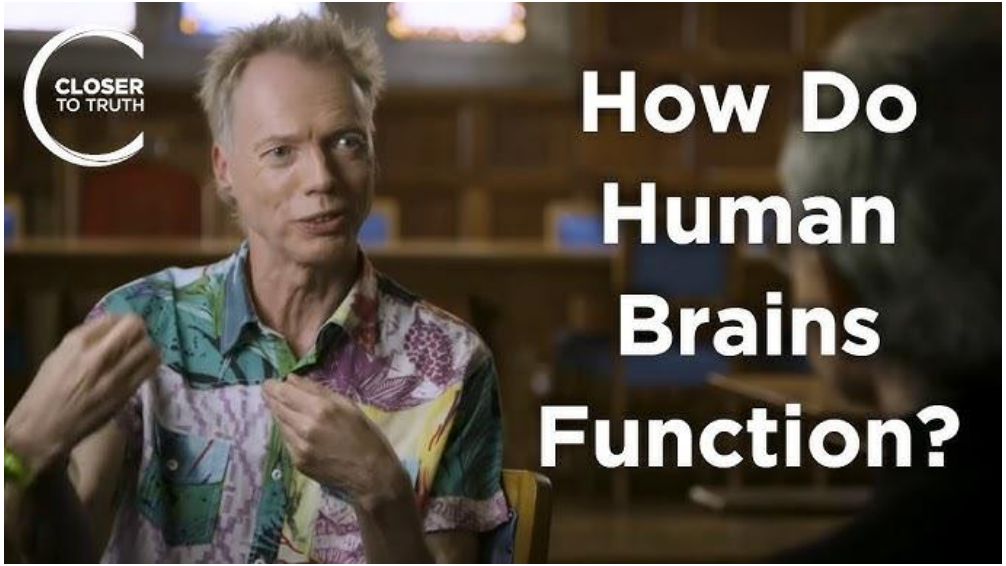
Second, the thought experiment does not show that *computers* can’t think. As we saw from Turing, since we human beings can carry out computations, we are computers, and we can think

Finally, the Chinese Room does not show that only brains can think; it is intended to show that brains must be doing more than the “mere shuffling of formal symbols” and that any artificial thinking machine would have to duplicate the specific neurophysiological processes responsible for thought in the human case


INPUT
牛对
DATA
RULE LEDGER
OUTPUT


SHORTLY BEFORE THE DEATH OF JOHN SEARLE:
I HAVE A CONFESSION. I'M ACTUALLY A CHINESE ROOM.


That whole system understands Chinese.







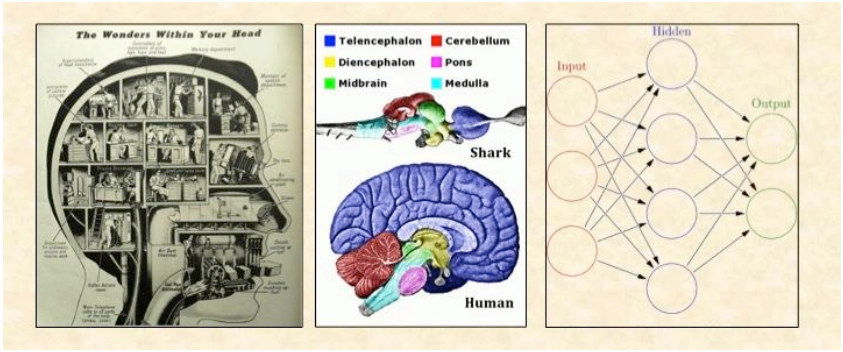
Suppose, for instance, that we **move beyond high-level symbolic representations**. If we implement finer, more complex structures (what Clark calls **micro-functionalism**), it's plausible that real understanding could emerge


ANY
QUESTIONS?

Thanks

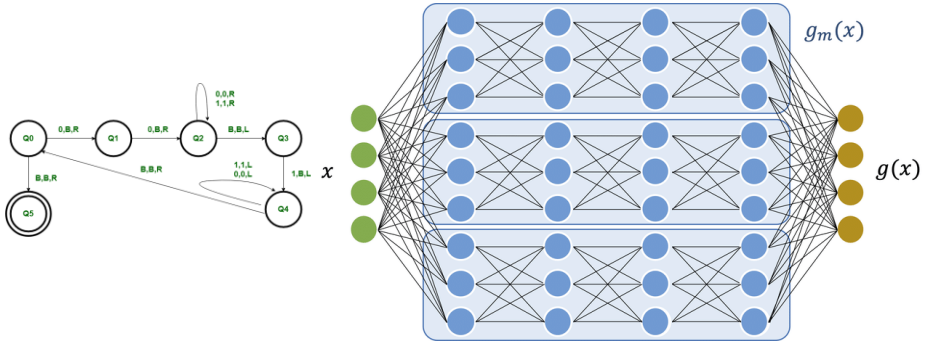


Connectionism: A Brain model of the Mind



Network models have a **neuron-like structures** to perform computations, **much like the brain**

GOFAI and classical AI models rely on serial processing, handling one operation at a time. The human brain, in contrast, **works in parallel**



intelligence emerges from the **interactions of many simple processing units**. As Rumelhart et al. (1986) put it, "The brain exemplifies the fact that it is all right to be very stupid if you're well connected."

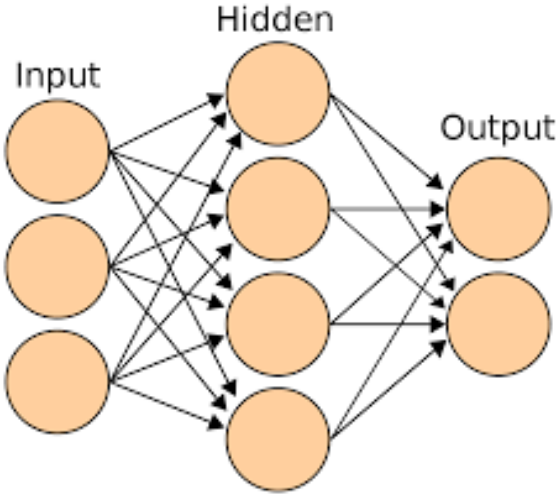


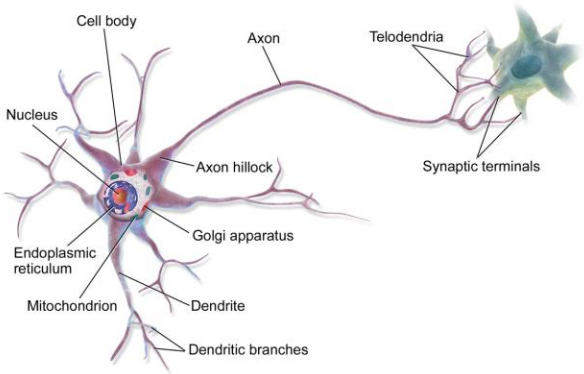


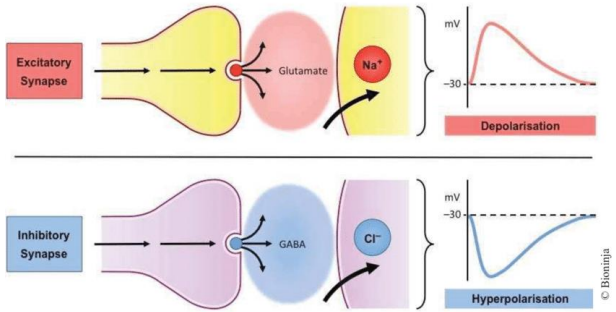
Connectionist models are built from **networks of neuron-like units**, each connected to others.

These units can fire at different levels, much like biological neurons. Units are treated as "on" if their activity level **crosses a certain threshold** and "off" if they don't.

Connections between these units can be either **excitatory** or **inhibitory**, just like neurons.



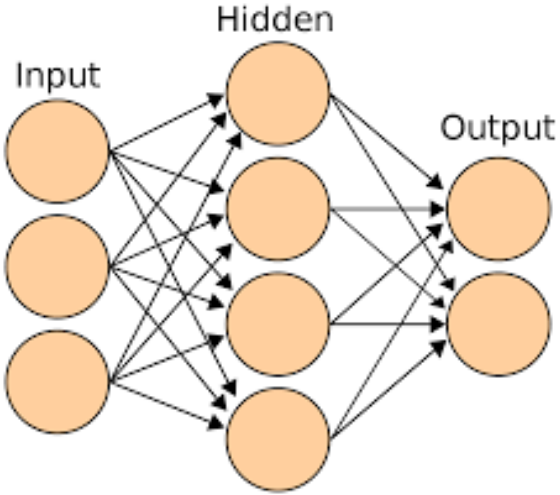


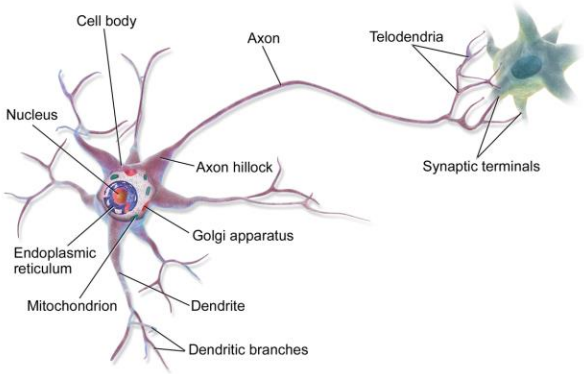


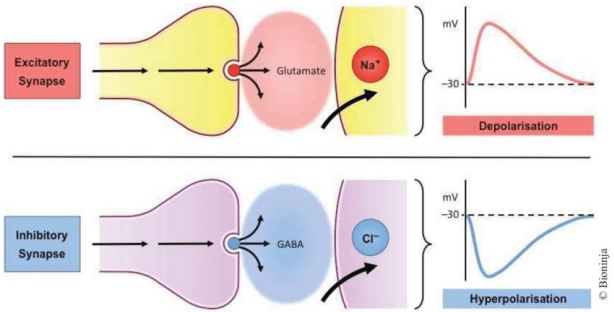
The hidden layer |where the system **processes information**) takes inputs from the previous layer (like sensory data or other signals) and transforms them in a way that helps the model learn patterns.

Each "node" (or artificial neuron) in the hidden layer combines and adjusts the input signals **using weights**, and then passes this processed information to the next layer.

The hidden layer's job is **to find connections and features in the data that aren't obvious in the raw input**, helping the model make better predictions or decisions

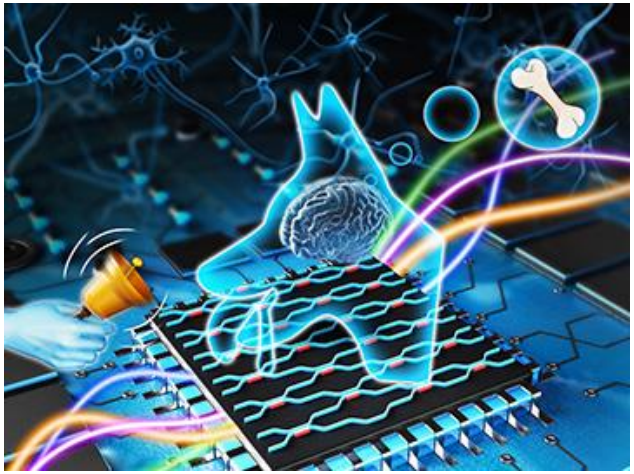


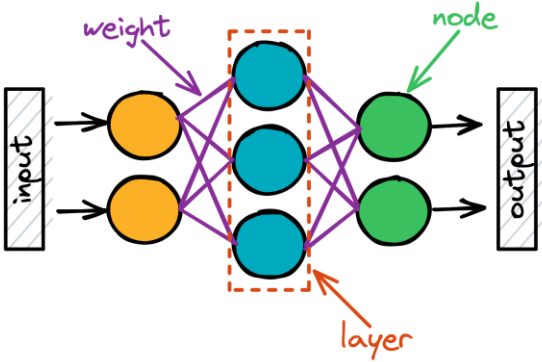


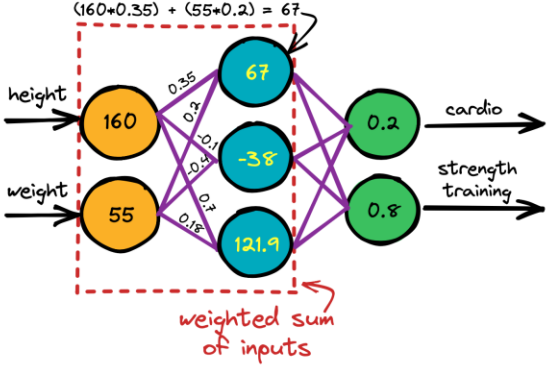


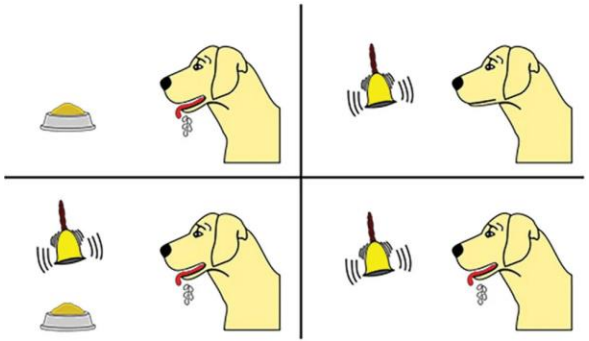
Learning in Connectionist Models

A defining feature of connectionist systems is their **ability to learn** by adjusting the **weights** of the connections between units. This weight-changing process is what enables the system to adapt based on experience. The weights determine **how much influence one unit has on another**, and through learning, the system adjusts these weights to improve performance.



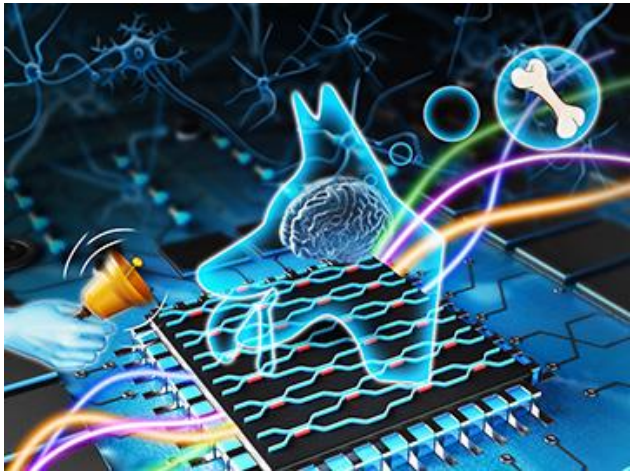


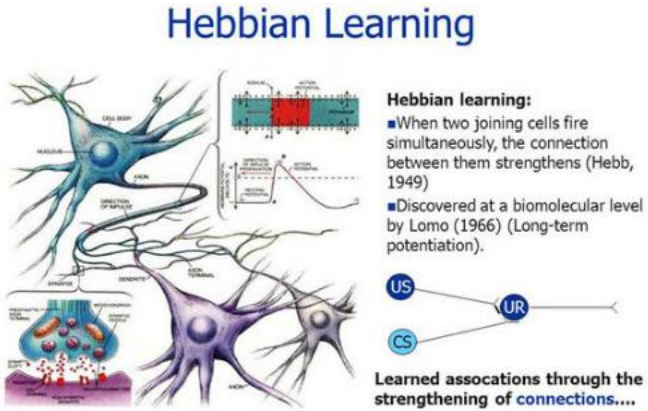


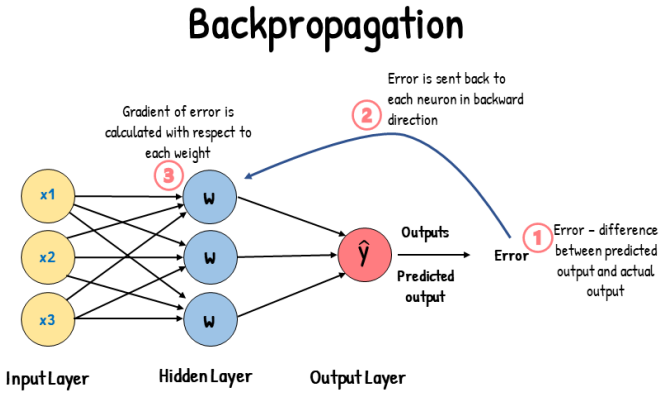


Learning in Connectionist Models

A defining feature of connectionist systems is their **ability to learn** by adjusting the **weights** of the connections between units. This weight-changing process is what enables the system to adapt based on experience. The weights determine **how much influence one unit has on another**, and through learning, the system adjusts these weights to improve performance.





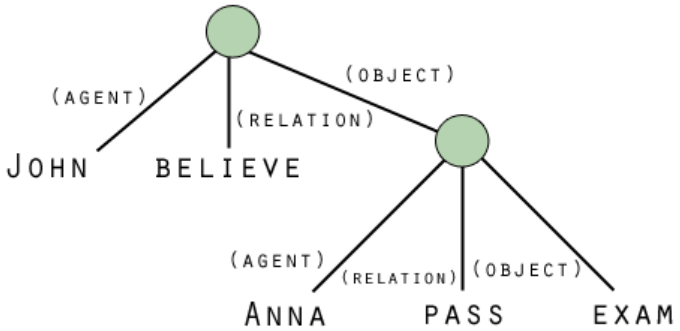


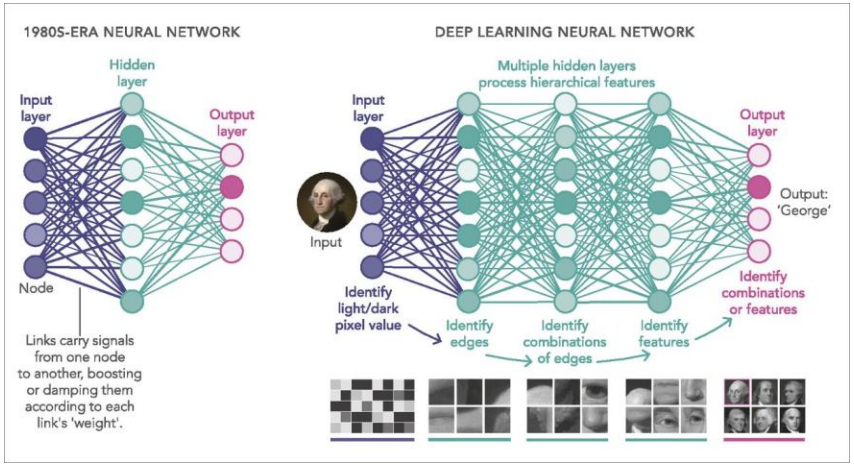
Representation in Connectionist Systems

Unlike GOFAI, where knowledge is stored in specific "boxes" or locations (e.g., beliefs, desires), connectionist models rely on **distributed representation**.

Instead of being stored in a single place, information is **spread across many units in the network**, making it more flexible and efficient.



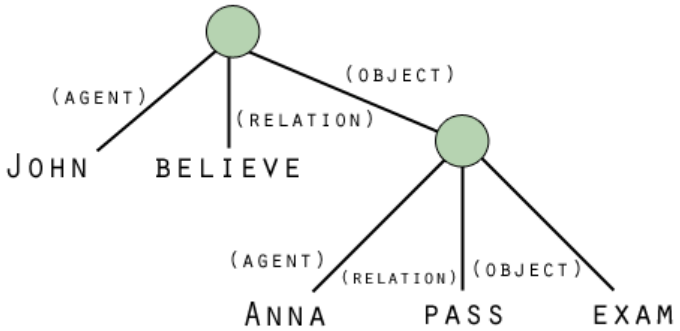


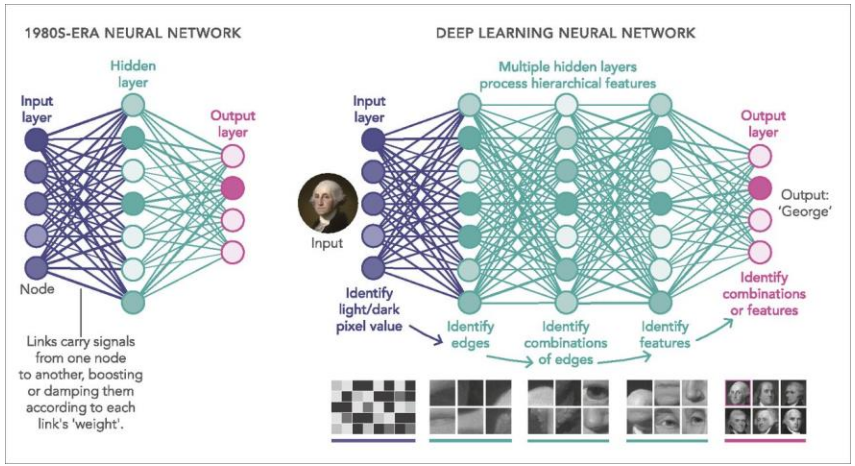


Damage Resistance and Graceful Degradation

Because information is distributed across the network, the system can cope with **incomplete or noisy data**. **Even if part of the system is damaged, it doesn't crash**; instead, performance degrades gradually, as is the case in biological brains, where neurons die regularly without catastrophic effects on cognition.

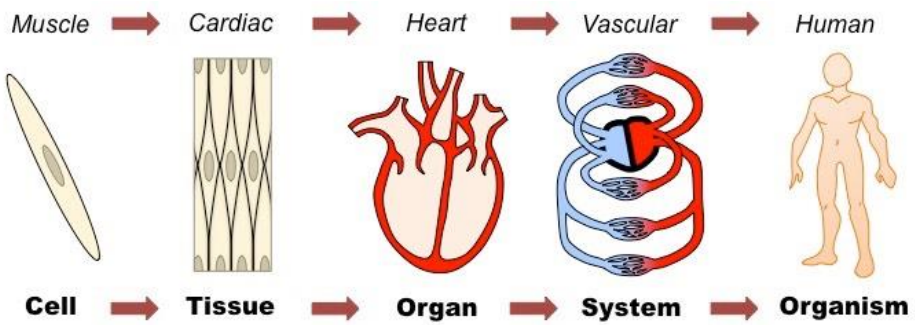






Emergent Properties in Connectionist Models

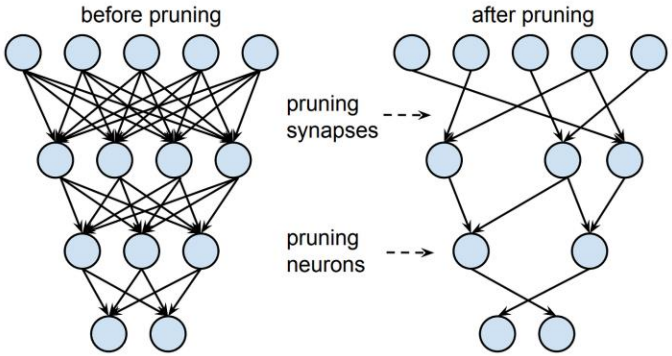
Emergent properties—characteristics that arise from the system as a whole but are not present in the individual units.





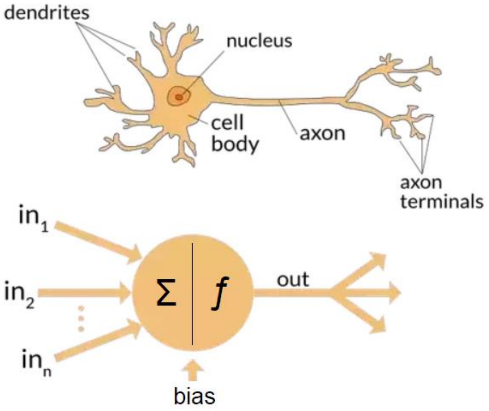
For instance, in a well-designed connectionist network, properties like **fault tolerance** and **pattern completion** emerge from the interactions between units, even though these abilities do not exist at the level of a single neuron-like unit.

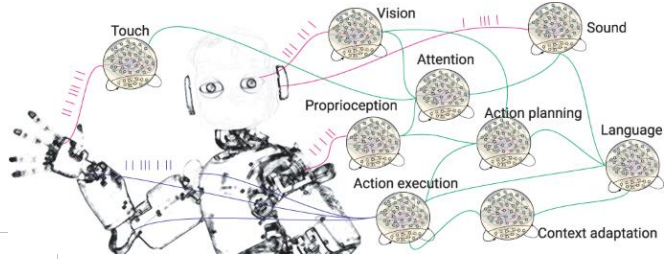


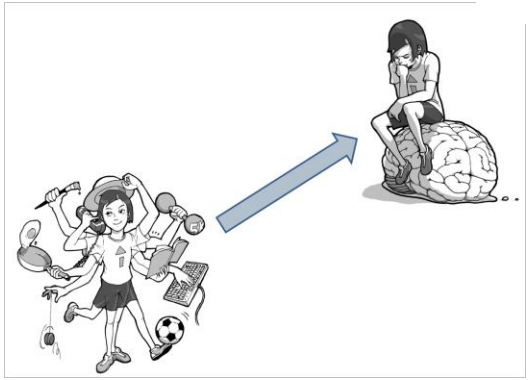

before pruning
after pruning
pruning synapses
pruning neurons



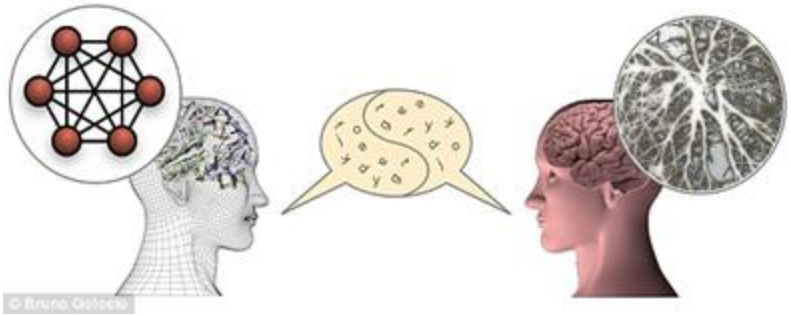

BIOLOGICAL
PLAUSIBILITY
EPIDEMIOLOGIST
STRICTEST
VACCINATION
BIOCHEMICAL
PUTATIVE
CENTURY
CASE
DISEASE
EVENT
DEBATES
BIOLOGIC
IMPROBABLE
CONTRACTED
COHERENT
VACCINE
SUFFICIENTLY
OUTCOME
RANDOMIZED
CRITERIA
THERAPY
REFUTE
SPECIFICITY
EMIGRANT
REFUTATION
STEERAGE
CAUSAL
EPIDEMIOLOGY
ASCRIBE
ABSURD
DEFINITIVE
MEDICINE
ESSAYIST
EXPLICABLE
EVALUATIONS
CONFER
TYPHUS
DATA
RUBELLA
FALLACY
FACTOR
CHIMNEY
ADVERSE
VERMIN
CLINICAL
PROCESS
SWEEPS
KNOWLEDGE
PLAUSIBLE
BIOLOGICAL
PLAUSIBILITY
RELATIONSHIP
BIOMEDICINE
DRUG
MERIT
SEMINAL
BIOLOGICALLY
SUPPORT
CAUSALITY
PROPOSED
BIOLOGY
CAUSATION
NEARER
TEMPORAL
EPIDEMIOLOGICAL
Adobe Stock | #48003514


dendrites
nucleus
cell body
axon
axon terminals
in1
in2
inn
Σ
f
out
bias


Touch
Vision
Sound
Attention
Proprioception
Action planning
Language
Action execution
Context adaptation





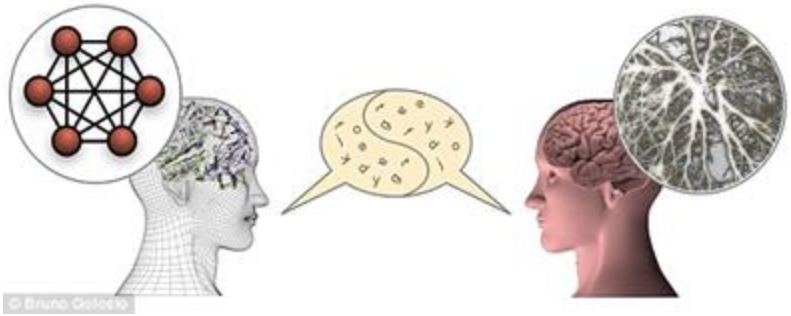


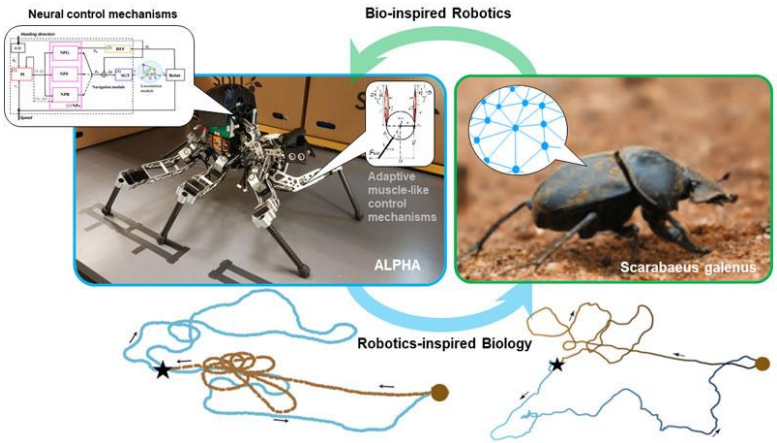
Are these systems **truly intelligent**?

Is this intelligence fundamentally **different from ours**?

How do we tell?






Neural control mechanisms
Bio-inspired Robotics
Adaptive muscle-like control mechanisms
ALPHA
Scarabaeus galenus
Robotics-inspired Biology









Classical Computational frameworks provide a powerful start for understanding the mind, but they potentially falls short when we consider **real-world fluency** and **the intricacies of embodied thought.**

Cognitive science is moving toward viewing the mind as ***fine-grained, multipurpose, and embodied*** – not a centralized, abstract system, but a fluid, distributed one, constantly shaped by interactions with the world.

As we integrate insights from **connectionism** and embodied cognition, we edge closer to understanding human cognition as *fine-grained, multipurpose, and deeply embodied*.

Thanks

