What is “artificial intelligence”?

- **Dartmouth Proposal** (McCarthy et al. 1955)

What is “artificial intelligence”?

Origin of “AI”!

A Proposal for the

DARTMOUTH SUMMER RESEARCH PROJECT ON ARTIFICIAL INTELLIGENCE

We propose that a 2 month, 10 man study of artificial intelligence be carried out during the summer of 1956 at Dartmouth College in Hanover, New Hampshire. The study is to proceed on the basis of the conjecture that every aspect of learning or any other feature of intelligence can in principle be so precisely described that a machine can be made to simulate it. An attempt will be made to find how to make machines use language, form abstractions and concepts, solve kinds of problems now reserved for humans, and improve themselves. We

What is “artificial intelligence”?

- **Dartmouth Proposal** (McCarthy et al. 1955): “The study is to proceed on the basis of the conjecture that every aspect of learning or any other feature of intelligence can in principle be so precisely described that a machine can be made to simulate it.”
- **John McCarthy** (2007): “It is the science and engineering of making intelligent machines, especially intelligent computer programs.” “The ultimate effort is to make computer programs that can solve problems and achieve goals in the world as well as humans.”
- **Wallace Marshall** (1987): “Artificial stupidity (AS) may be defined as the attempt by computer scientists to create computer programs capable of causing problems of a type normally associated with human thought.”

Okay, so what is “intelligence”? (How do we know if we've succeeded?)

What do you think? What is the defining feature of “intelligence”?

- A capacity to solve problems.
- Ability to use knowledge to devise novel solutions to problems.
- Ability to learn something and improve upon it; ability to make a decision.
- Two features: reasoning, adaptability.

Okay, so what is “intelligence”? (How do we know if we've succeeded?)

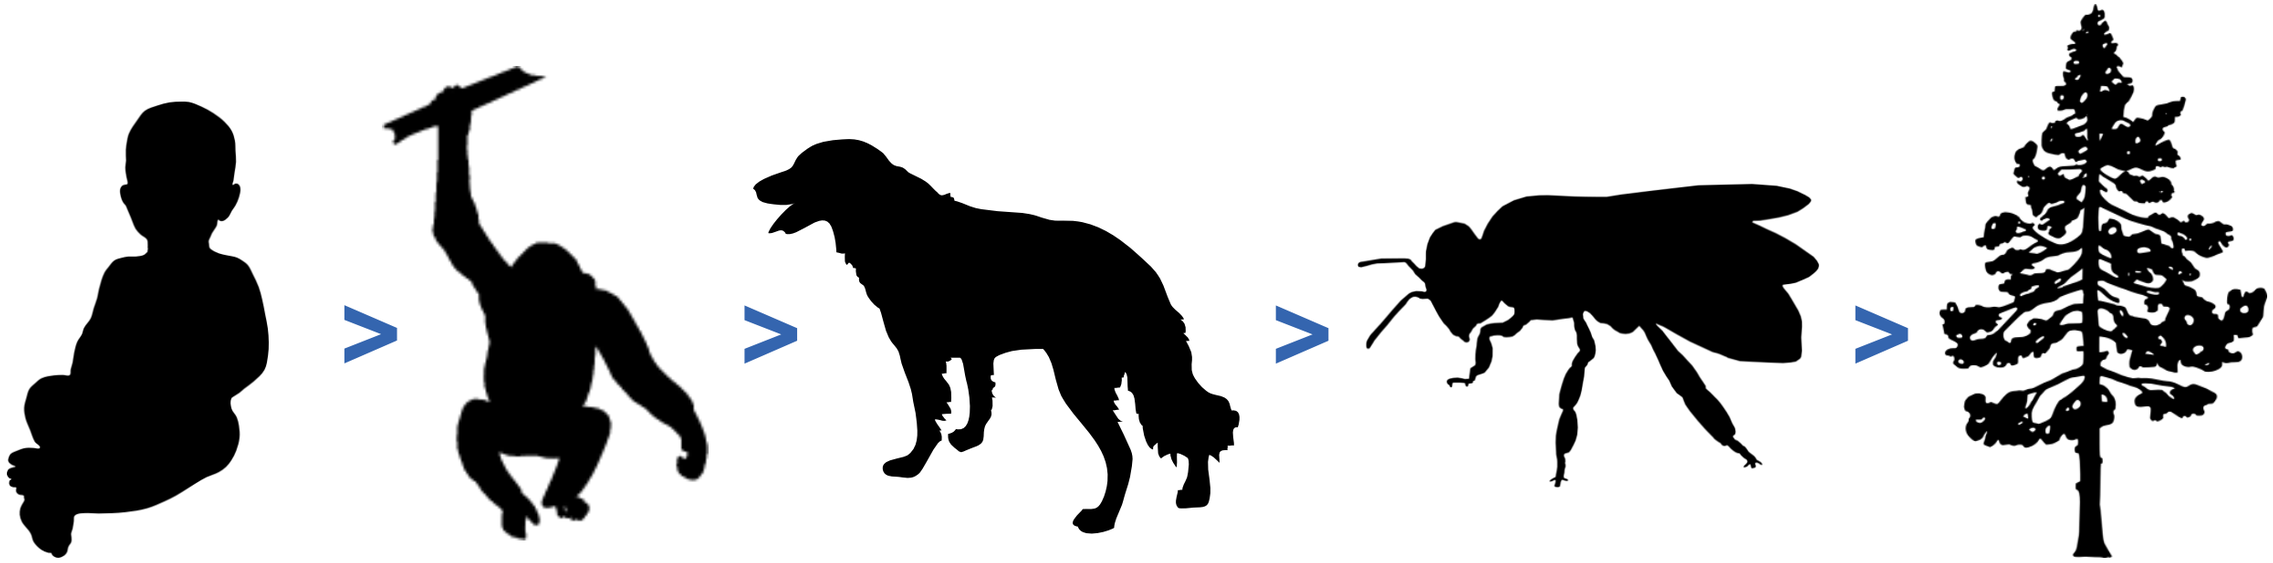

Okay, so what is “intelligence”? (How do we know if we've succeeded?)

- **McCarthy** (2007): “Intelligence is the computational part of the ability to achieve goals in the world. Varying kinds and degrees of intelligence occur in people, many animals and some machines.”
- **Allen Newell** (1990): An intelligent system “operates in real-time; exploits vast amounts of knowledge; tolerates erroneous, unexpected, and possibly unknown inputs; uses symbols and abstractions; communicates using some form of natural language; learns from the environment; and exhibits adaptive goal-oriented behavior.”
- **Kurzweil** (1990): “In summary, there appears to be no *simple* definition of intelligence that is satisfactory to most observers, and most would-be definers of intelligence end up with long checklists of its attributes.”
- **Alan Turing** (1950): The “imitation game” (now called the “Turing test”).

Turing, AM (1950) Computing Machinery and Intelligence. *Mind*: 54.

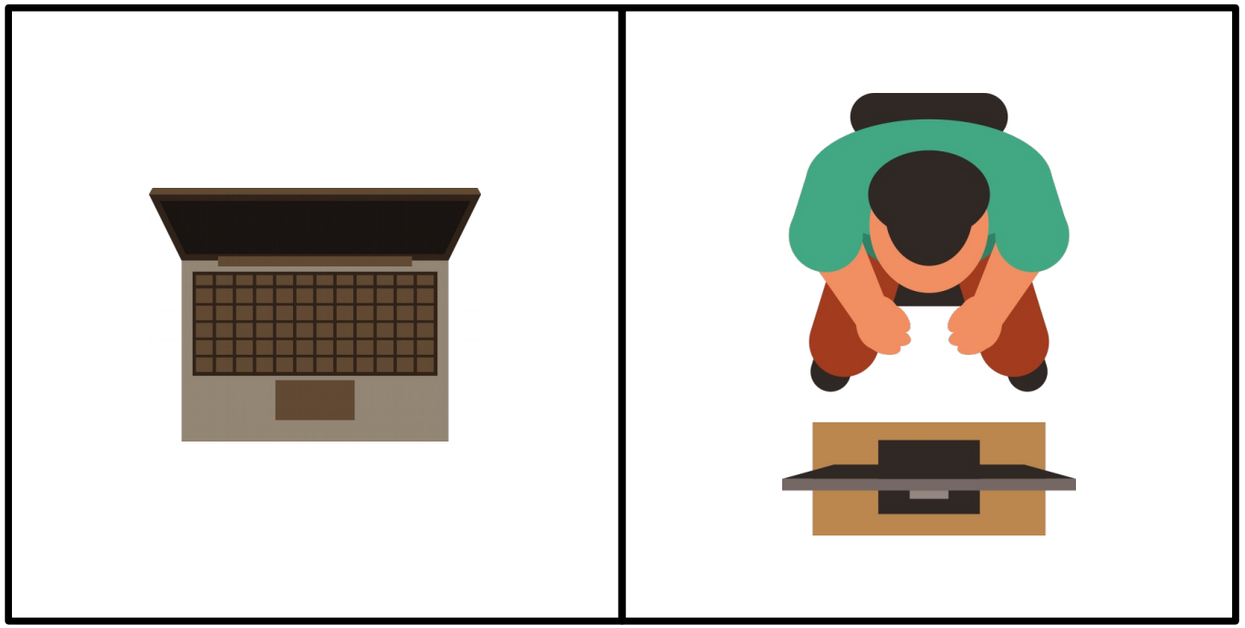

Which room has the computer?

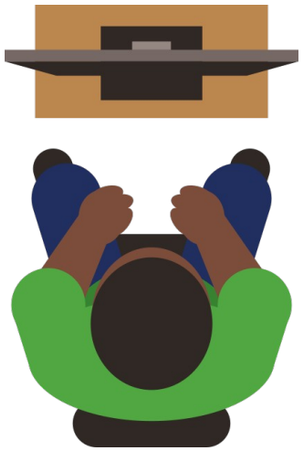

Sufficient condition for demonstrating intelligence*, but not *necessary*!

* Not everyone agrees with this. See, e.g., Searle (1980).

Whirlwind tour of the history of, and principle approaches to, AI research

Pre-1956: Early foundations

Ada Lovelace

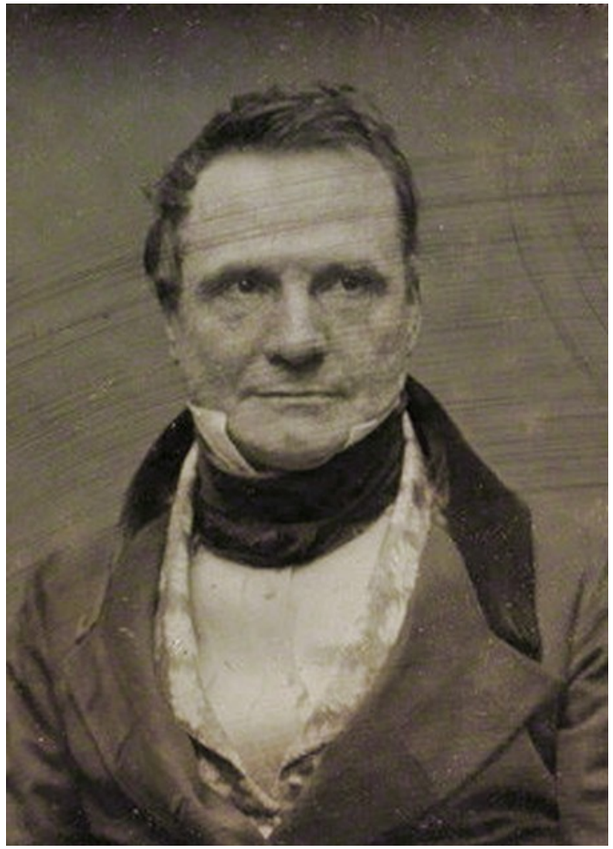

Charles Babbage

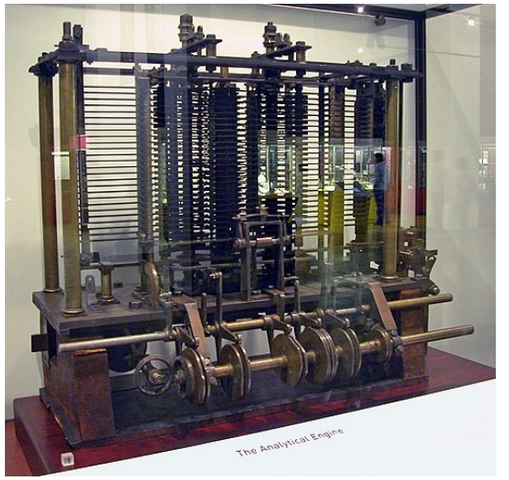

Babbage's
“Analytical Engine”

Bruno Barral (CC-SA)

Lovelace (1843): “Again, it might act upon other things besides number…”
“Supposing, for instance, that the fundamental relations of pitched sounds … were susceptible of such expression and adaptations, the engine might compose elaborate and scientific pieces of music of any degree of complexity or extent.”

1950 1960 1970 1980 1990 2000 2010 2020

Pre-1956: Early foundations

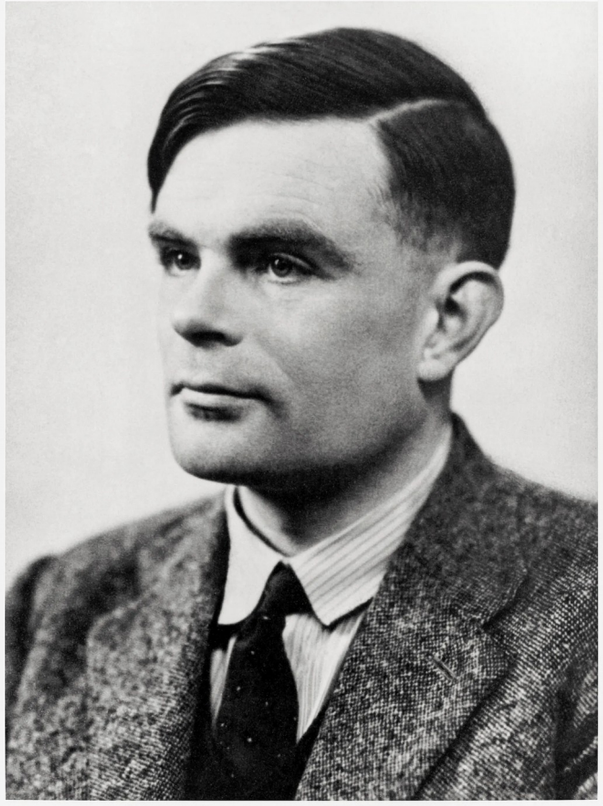

The New Yorker/Alamy

Alan Turing

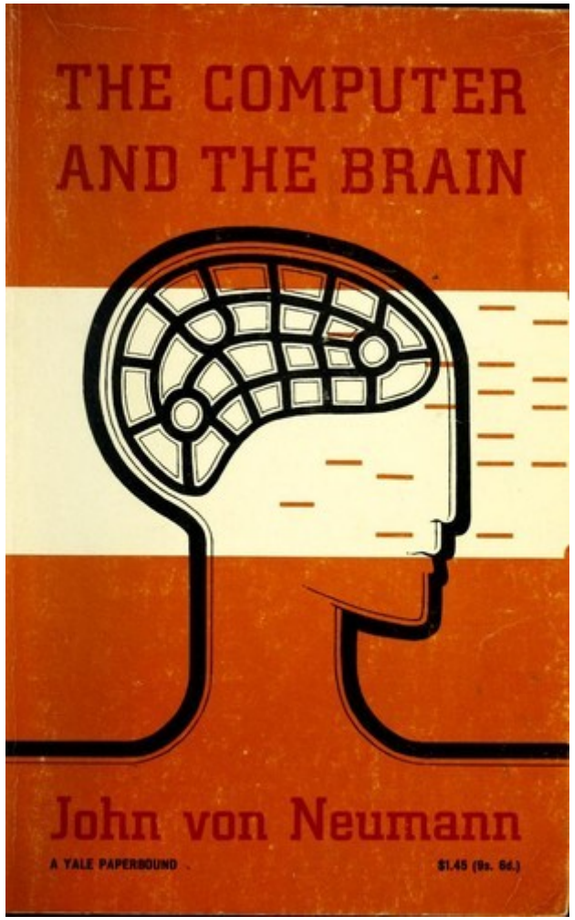

John von Neumann (1956)

An excellent read.
Short, accessible.

1950 1960 1970 1980 1990 2000 2010 2020

A PROPOSAL FOR THE

DARTMOUTH SUMMER RESEARCH PROJECT

ON ARTIFICIAL INTELLIGENCE

J. McCarthy, Dartmouth College
M. L. Minsky, Harvard University
N. Rochester, I. B. M. Corporation
C. E. Shannon, Bell Telephone Laboratories

August 31, 1955

1956 - mid 1970s: First AI boom

Incredibly, the major computational approaches to AI were all established during this time.

We might think of them in two categories (both of which we'll cover in this class):

- Top-down methods
- Bottom-up methods

The top-down approach to AI (symbolic AI)

- Focus on knowledge, symbolic (abstract) manipulation, reasoning, and problem-solving algorithms.
- AI achieved by direct engineering.
- Example: Common-sense reasoning about objects.

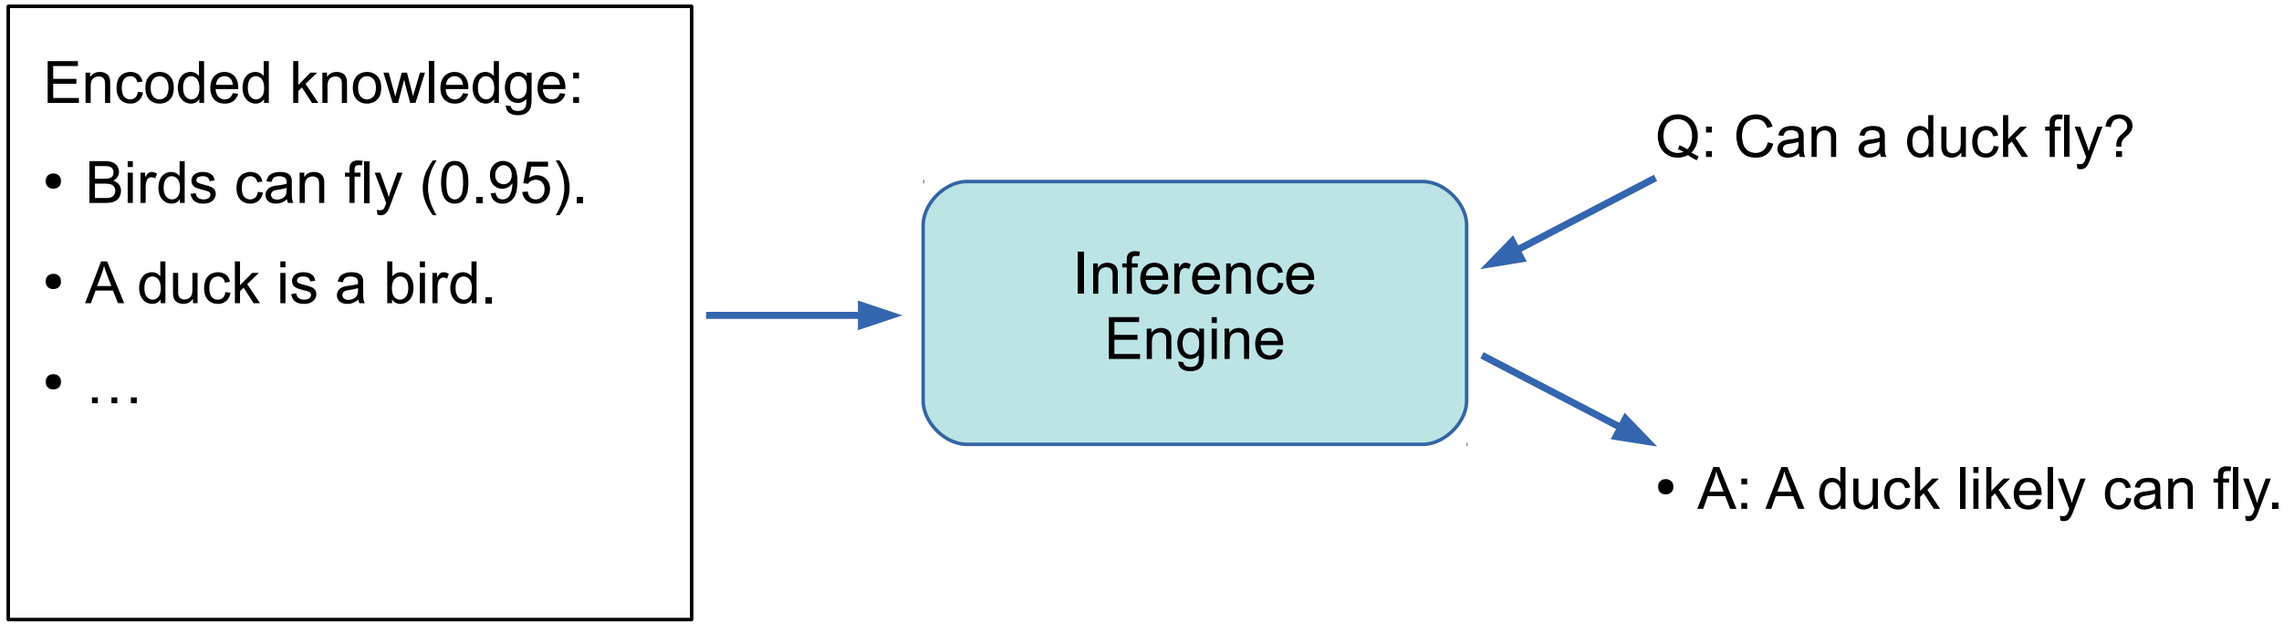

The top-down approach to AI (symbolic AI)

- Focus on knowledge, symbolic (abstract) manipulation, reasoning, and problem-solving algorithms.
- AI achieved by direct engineering.
- Example: Playing checkers.

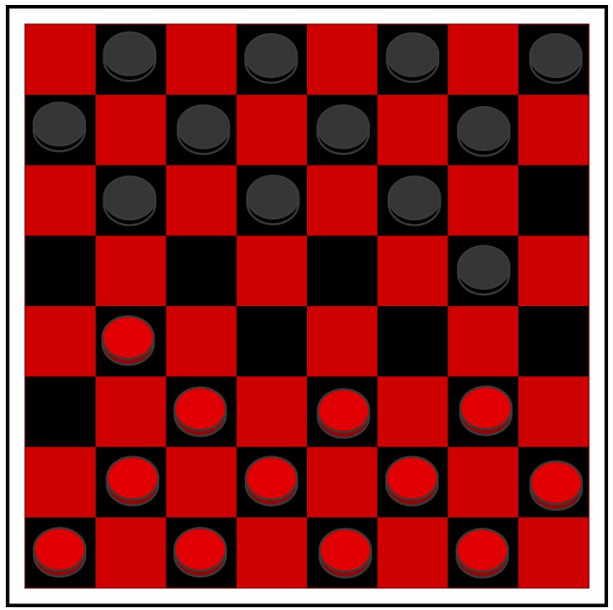

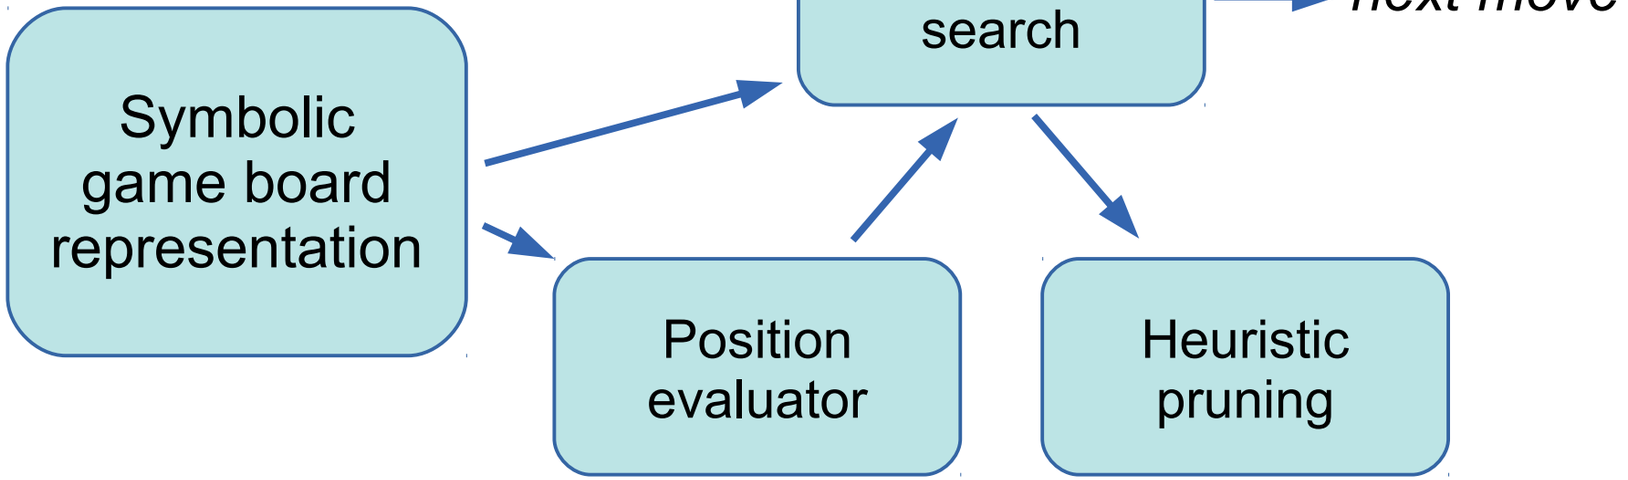

The bottom-up approach to AI (machine learning and friends)

- Focus on inputs (e.g., data, environment) and low-level methods of processing.
- AI achieved by learning, self-organization, or optimization.
- Example: Image recognition.

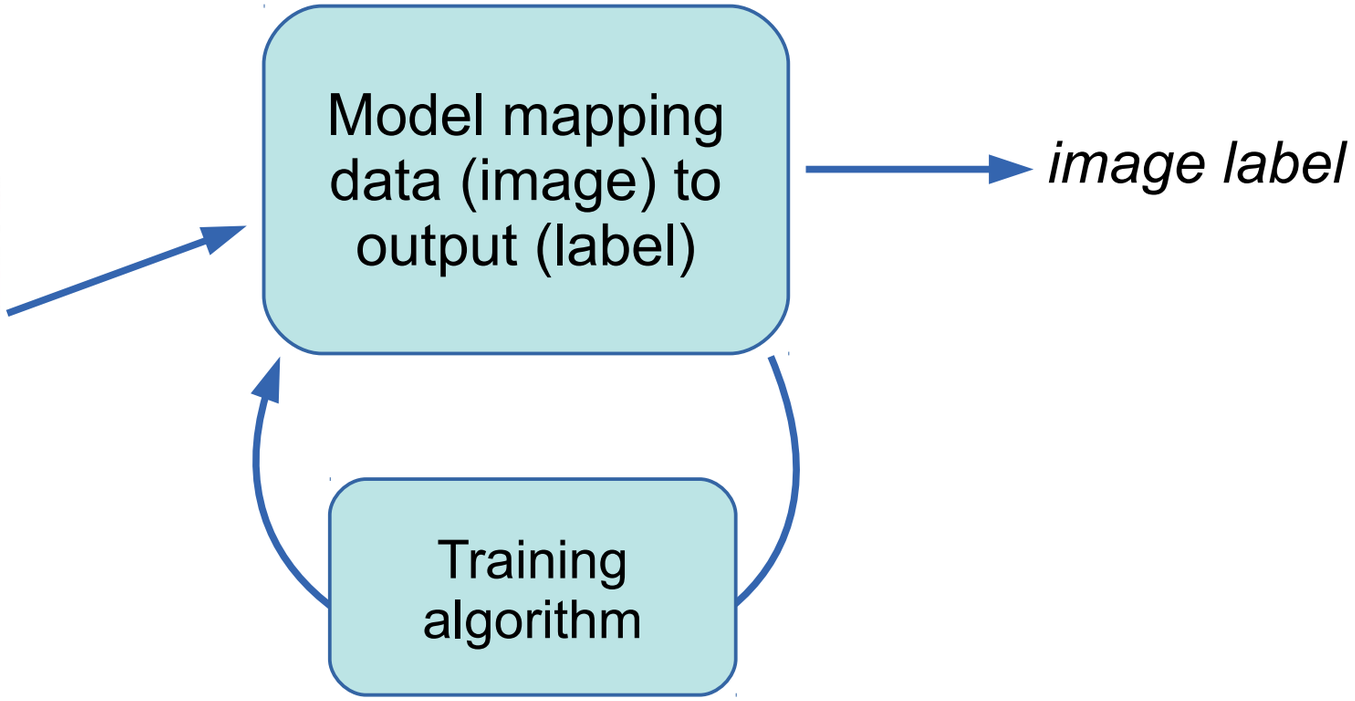

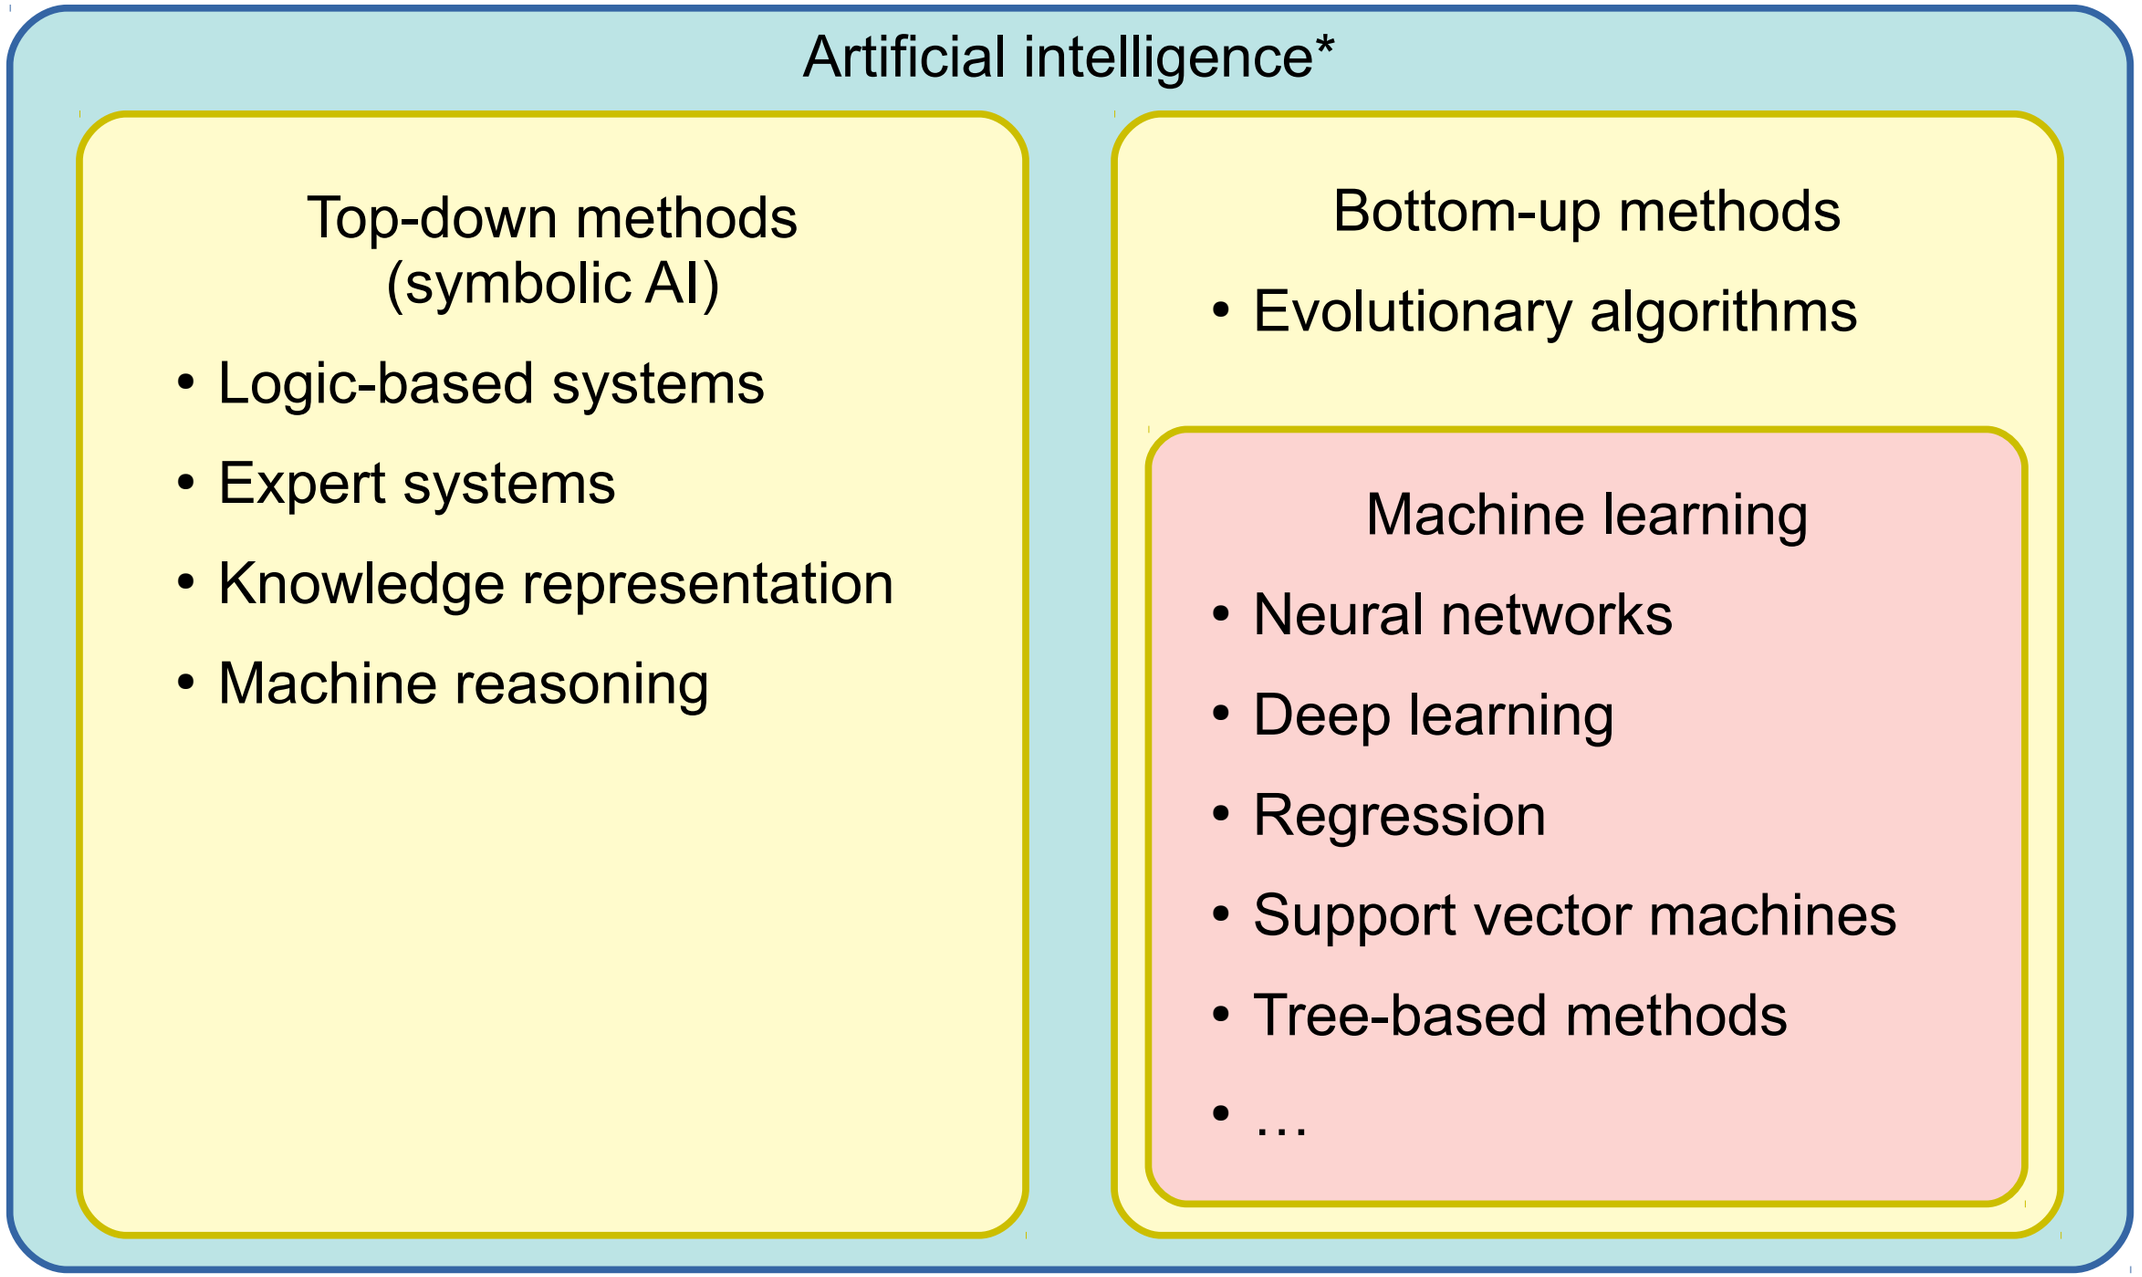

* I'm focusing on computing approaches and not including robotics, but it also plays a major role in AI.

1956 - mid 1970s: First AI boom

Incredible progress (game playing, theorem proving, image recognition) led to incredible optimism.

- **Herbert Simon** (1960): “Technologically, as I have argued earlier, machines will be capable, within twenty years, of doing any work that a man can do.”
- **Kurzweil** (1999): “Perhaps Simon was intending a favorable comment on the capabilities of women...”
- **Marvin Minsky** (1967): “Within a generation, I am convinced, few compartments of intellect will remain outside the machine's realm – the problems of creating "artificial intelligence" will be substantially solved.”

mid 1970s - early 1980s: First AI winter

- **James Lighthill's report** (1973): “Most workers in AI research and in related fields confess to a pronounced feeling of disappointment in what has been achieved in the past twenty-five years. … In no part of the field have the discoveries made so far produced the major impact that was then promised.”
- **Hans Moravec** (1988): “… it has become clear that it is comparatively easy to make computers exhibit adult-level performance in solving problems on intelligence tests or playing checkers, and difficult or impossible to give them the skills of a one-year-old when it comes to perception and mobility.”
- *The hard things were easy, and the easy things were hard.*
- *The hard things were hard, and the easy things were <u>really</u> hard.*

early 1980s - early 1990s: Second AI boom (summer?)

- “Expert systems” become a multi-million-dollar industry.
- Propelled by AI that was *useful*.
- E.g., MYCIN did as well as, or better than, human specialists in diagnosing and prescribing treatments for bacterial disease (Yu et al. 1979).
- Emergence of a major AI hardware industry (Lisp machines).

What is an “expert system”?

- Symbolic AI system for answering questions and reasoning in a narrow domain of expertise.
- E.g.: medical diagnoses, CPU chip design, safety monitoring.

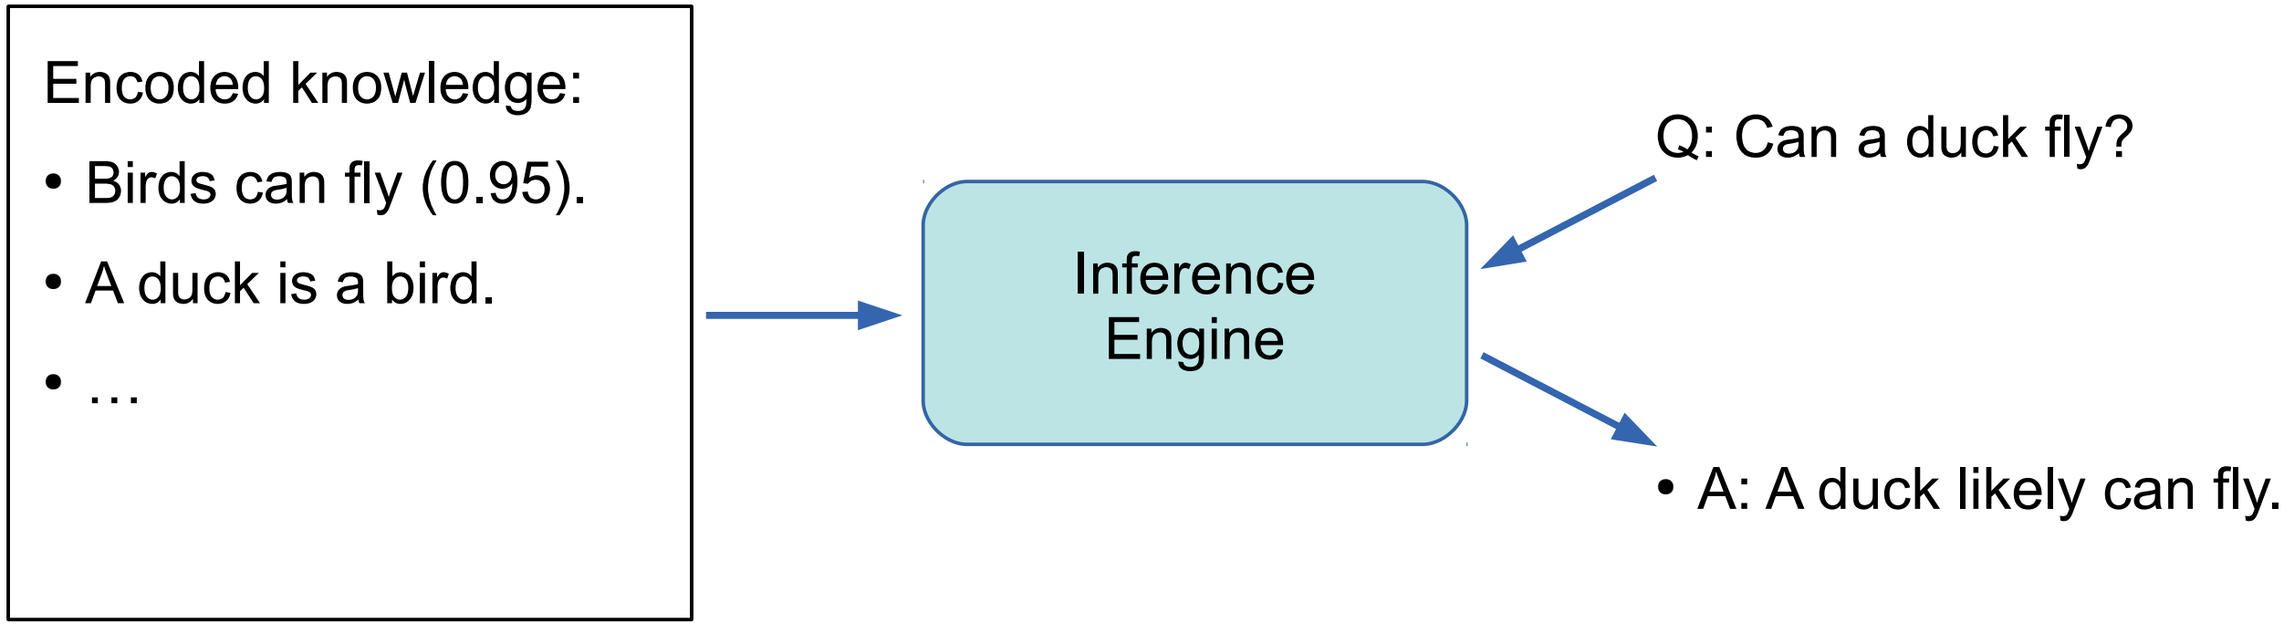

1990s: Second AI winter

- Expert systems and applications were over-hyped. Investors abandoned the industry.
- The AI hardware industry collapsed due to the rise of powerful PCs.

2000s – today: Third AI boom

- The rise of neural networks and deep learning.
- Key methods all developed decades earlier!
- Largely driven by:
  - Data (we have lots!)
  - Hardware (especially highly parallel GPUs)

What can we learn from the history of AI?

- AI researchers have been making steady progress since the 1950s.
- AI is seductive.
- AI winters were mostly about over-hyped commercial and government interest and expectations.
- AI winters weren't really about failures of AI. AI technologies, including many “failures” are relevant and widely useful today, *especially in science*.

So, are we riding another bubble that's about to collapse?

- Beats me. Some experts certainly think so.
- But, progress over the last decade has been real and ground-breaking.
- Of course, there have been some wild and crazy claims.
- **Sam Altman**, CEO of OpenAI (billions in investor funding) (2019):
  - “We’ve made a soft promise to investors that, ‘Once we build a generally intelligent system, that basically we will ask it to figure out a way to make an investment return for you.'”

How good are today's AI systems?

- Ground-breaking results over the last decade in a variety of applications.
- But, soccer balls and referees…
- Andrej Karpathy (Sr. Director of AI at Tesla):

Why is this funny?

"To me, examples like this illustrate that we are missing many crucial pieces of the puzzle…" "The road ahead is long, uncertain and unclear."

"We are really, really far away."

So, why are we still so far away from “solving” artificial intelligence?

- Perhaps we're not smart enough to build something as smart as we are.
  - Kurzweil (1990, 1999) argues that the intelligence of evolution is extraordinarily low, yet it was able to engineer creatures that far surpassed it in intelligence. We are more intelligent than evolution, *ergo* we should be able to do the same. (I am paraphrasing here.)
  - Right, but  .

    ©Cody Pope
  - Bottom line: We don't know what preconditions are necessary for intelligence to emerge.
- Perhaps our computers aren't powerful enough for general intelligence.

How much computing power does general intelligence require?

- How do we measure computing capacity?
  - By counting FLOPS.
  - A FLOP is a floating-point operation; e.g. 3.1416 + 2.7182.
  - FLOPS = floating-point operations per second.

What is the computing capacity of the human brain?

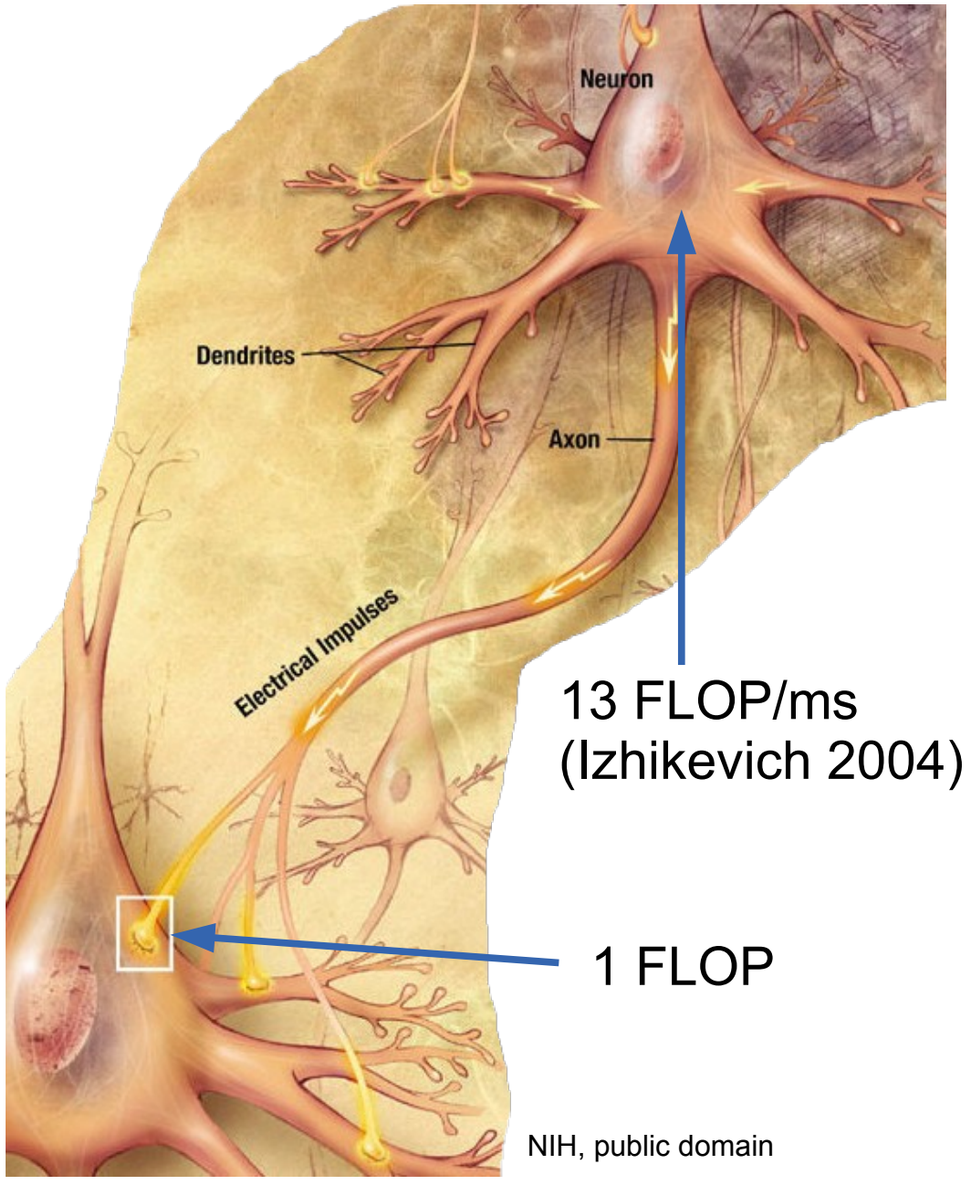

NIH, public domain

- 84.6 billion neurons in a human brain (Azevedo et al. 2009).
- ~7,000 synapses / neuron (Pakkenberg et al. 2003).
- Upper limit of ~2 action potentials (AP) / second / per neuron, on average (Lennie 2003).
- So, ~1.2×10^{15} FLOPS for AP accumulation
- and ~1.1×10^{15} FLOPS for firing decisions
- = ~2.3×10^{15} FLOPS for the human brain.

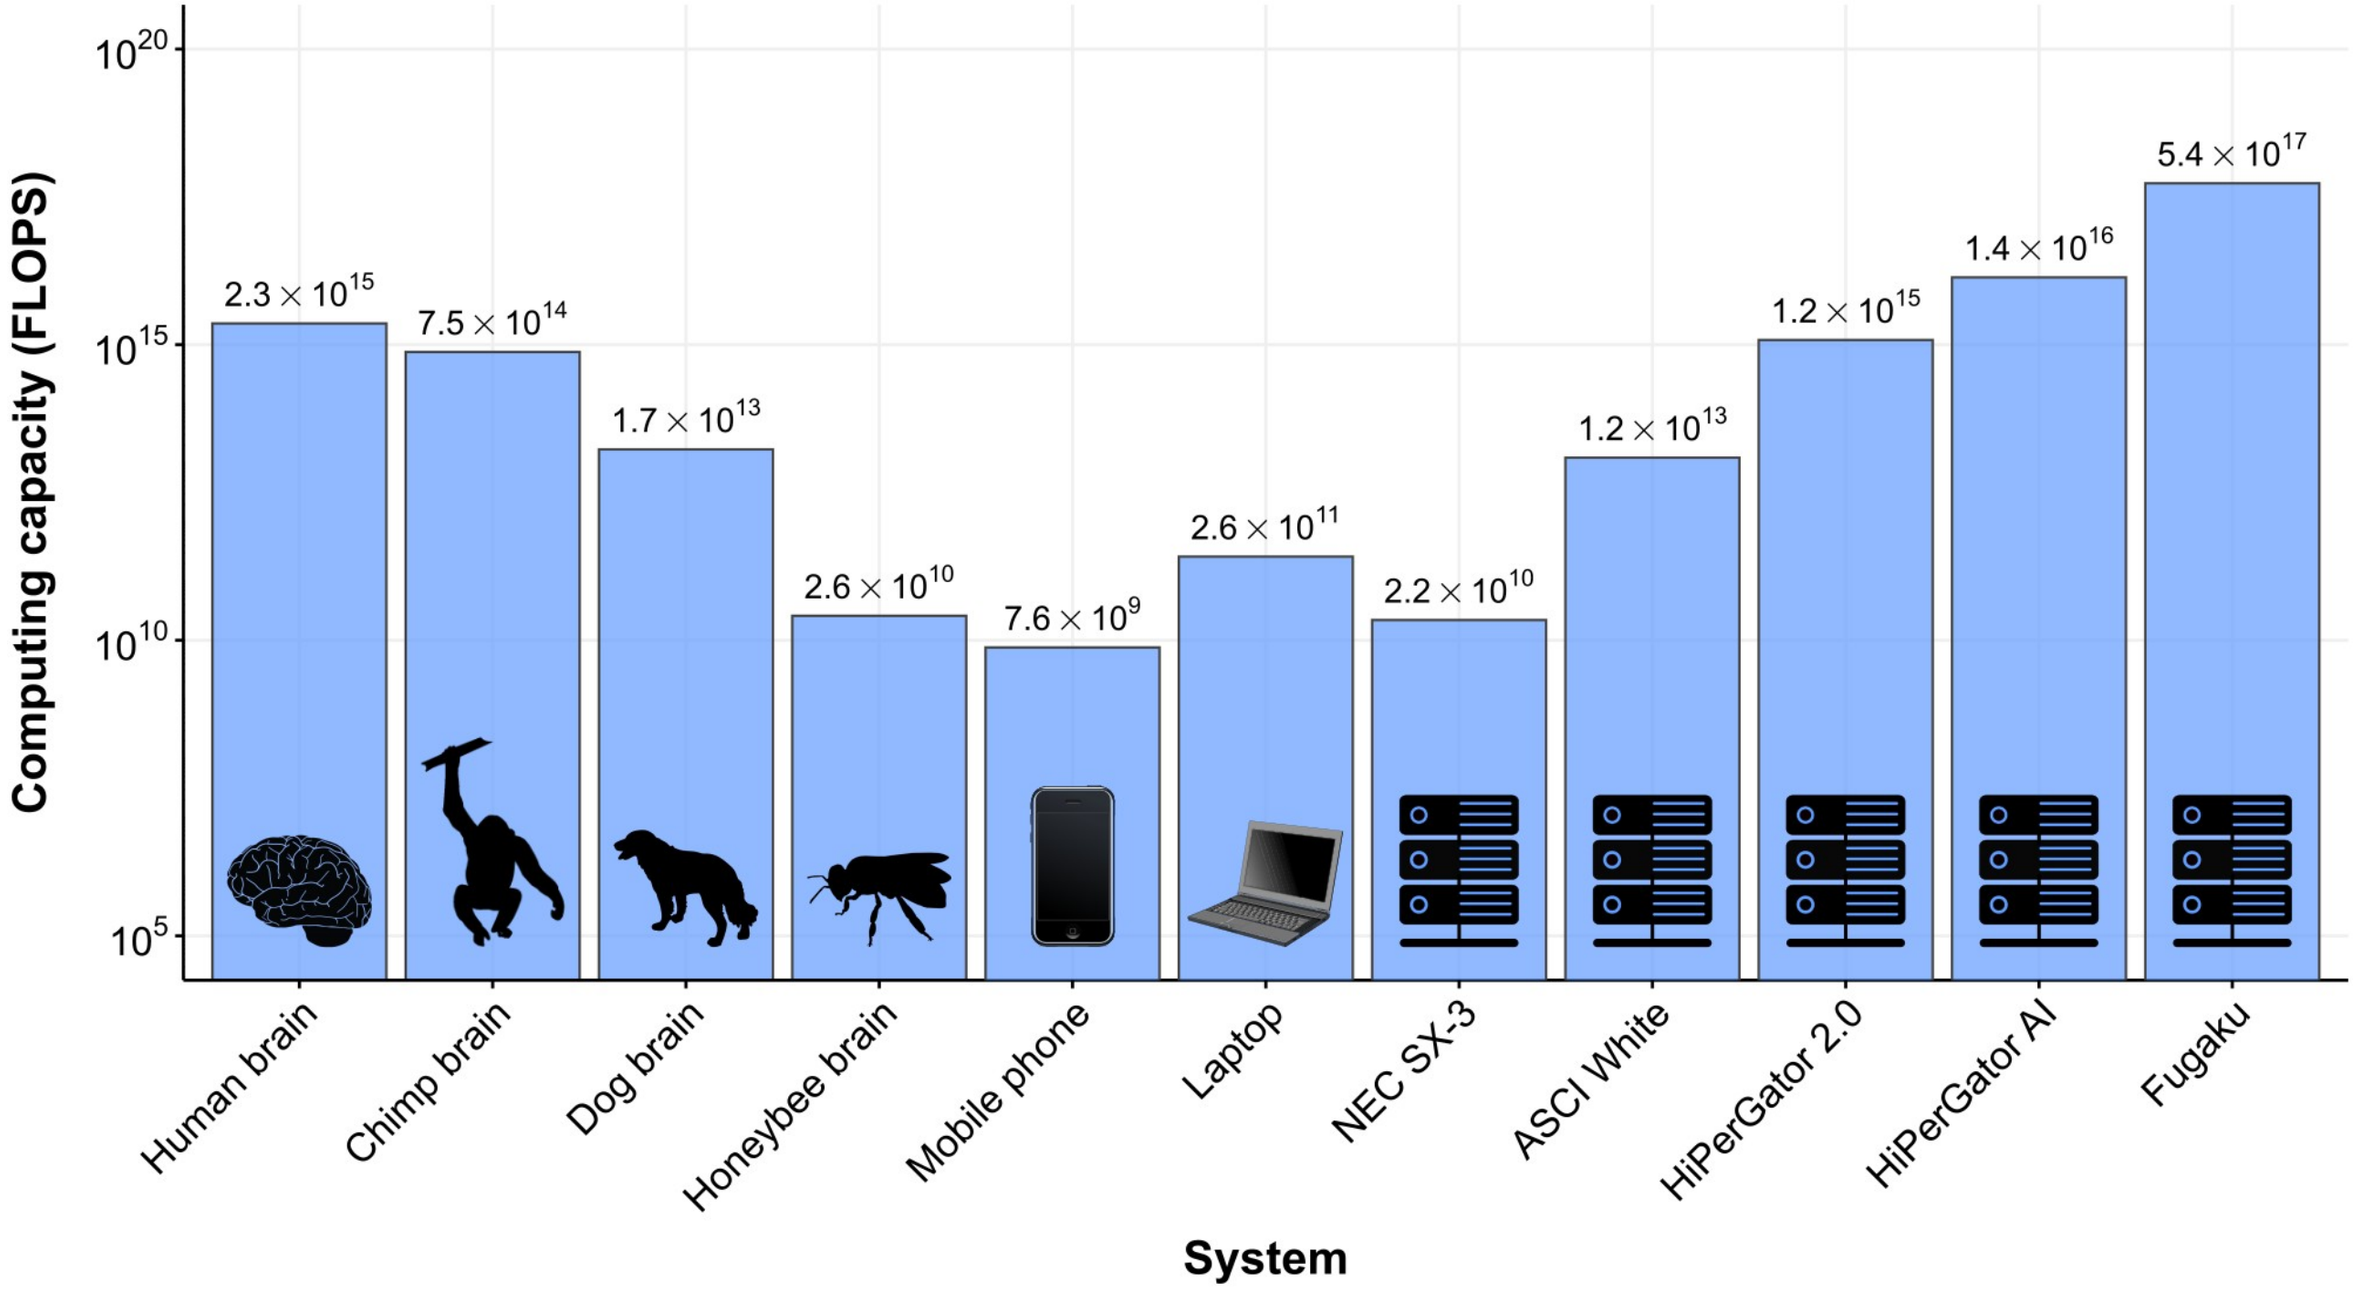
Computing capacity (FLOPS)
10^{20}
10^{15}
10^{10}
10^{5}
2.3×10^{15}
7.5×10^{14}
1.7×10^{13}
2.6×10^{10}
7.6×10^{9}
2.6×10^{11}
2.2×10^{10}
1.2×10^{13}
1.2×10^{15}
1.4×10^{16}
5.4×10^{17}
Human brain
Chimp brain
Dog brain
Honeybee brain
Mobile phone
Laptop
NEC SX-3
ASCI White
HiPerGator 2.0
HiPerGator AI
Fugaku
System

So, why are we still so far away from “solving” artificial intelligence?

- Perhaps we're not smart enough to build something as smart as we are. (*It's impossible.*)
- Perhaps our computers aren't powerful enough for general intelligence. (*It's a hardware problem.*)
- Pedro Domingos (2018): “It must be that we have a better learning algorithm in our heads than anything we’ve come up with for machines.” (*It's a software problem.*)

- *The field is wide open, and foundational questions remain largely unanswered!*
- *Countless applications to science remain unexplored.*