

Artificial Intelligence

What is Intelligence?

What is **Artificial** Intelligence?

John McCarthy (1927 - 2011)



Dartmouth Summer Research Project on Artificial Intelligence - 1956

- McCarthy himself, Allen Newell, Herbert Simon, Claude Shannon, and Marvin Minsky
- Groundwork for AI were first introduced



What is Artificial Intelligence?

"It is the science and engineering of making intelligent machines, especially intelligent computer programs. It is related to the similar task of using computers to understand human intelligence, but AI does not have to confine itself to methods that are biologically observable."

- John McCarthy, 2007

Yes, but what is intelligence?

"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."

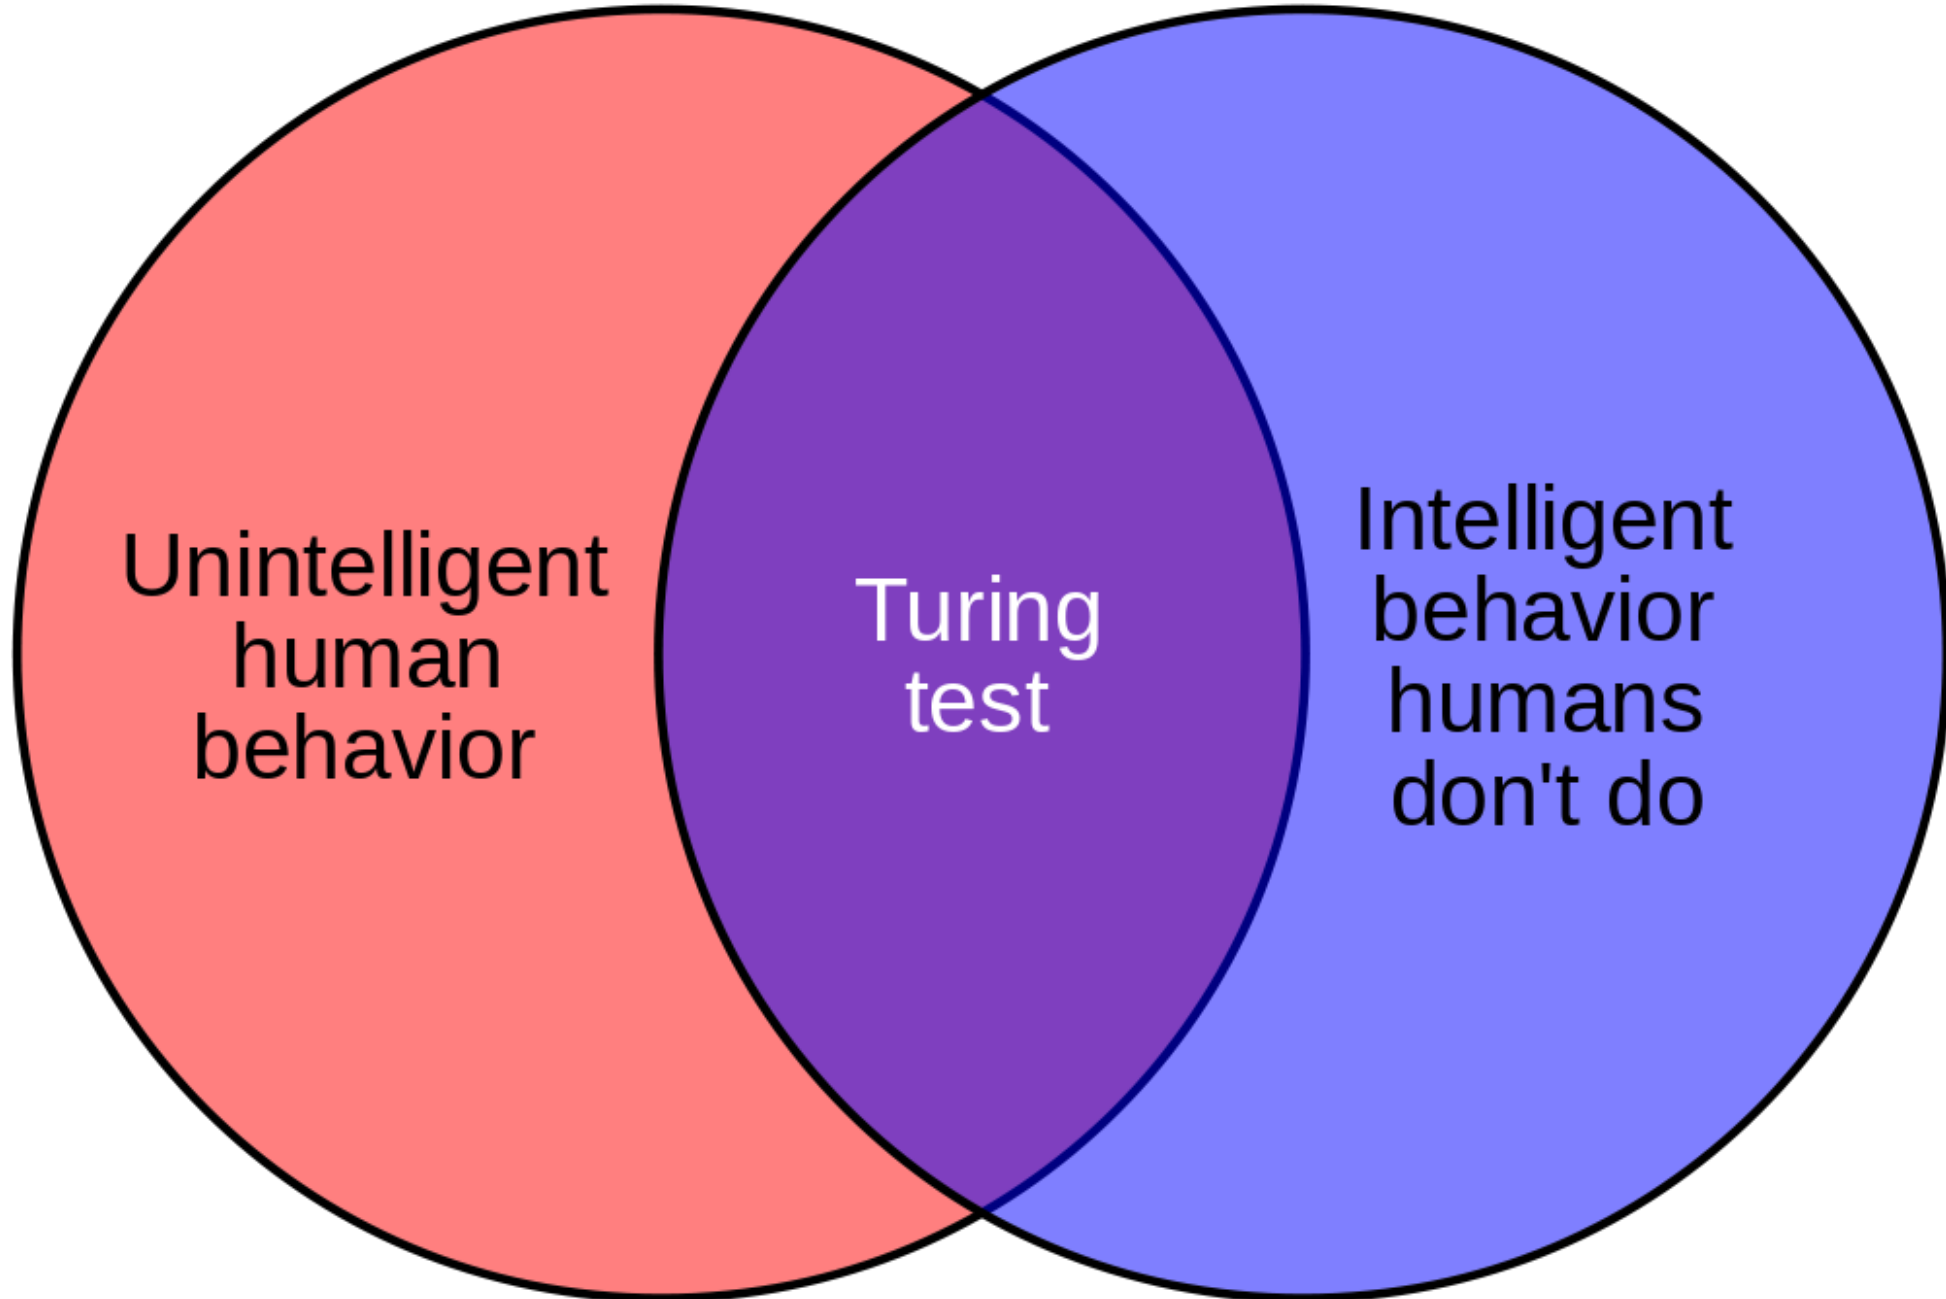
- John McCarthy, 2007

Isn't there a solid definition of intelligence that doesn't depend on relating it to human intelligence?

"Not yet. The problem is that we cannot yet characterize in general what kinds of computational procedures we want to call intelligent. We understand some of the mechanisms of intelligence and not others."

- John McCarthy, 2007

Human behavior Intelligent behavior



Alan Turing (1912 - 1954)

"I propose to consider the question, 'Can machines think?'"

- Turing, 1950



Turing Test



Chinese Room

- [John Searle](#) 1980: “Minds, Brains, and Programs”
- Instructions for input/out of Chinese characters
- Intelligent, or just good at following directions?



Strong AI vs. Weak AI

- **Strong AI** - matches human intelligence and is capable of performing *any* task
- **Weak AI** - only designed to perform a *specific subset* of intelligent actions

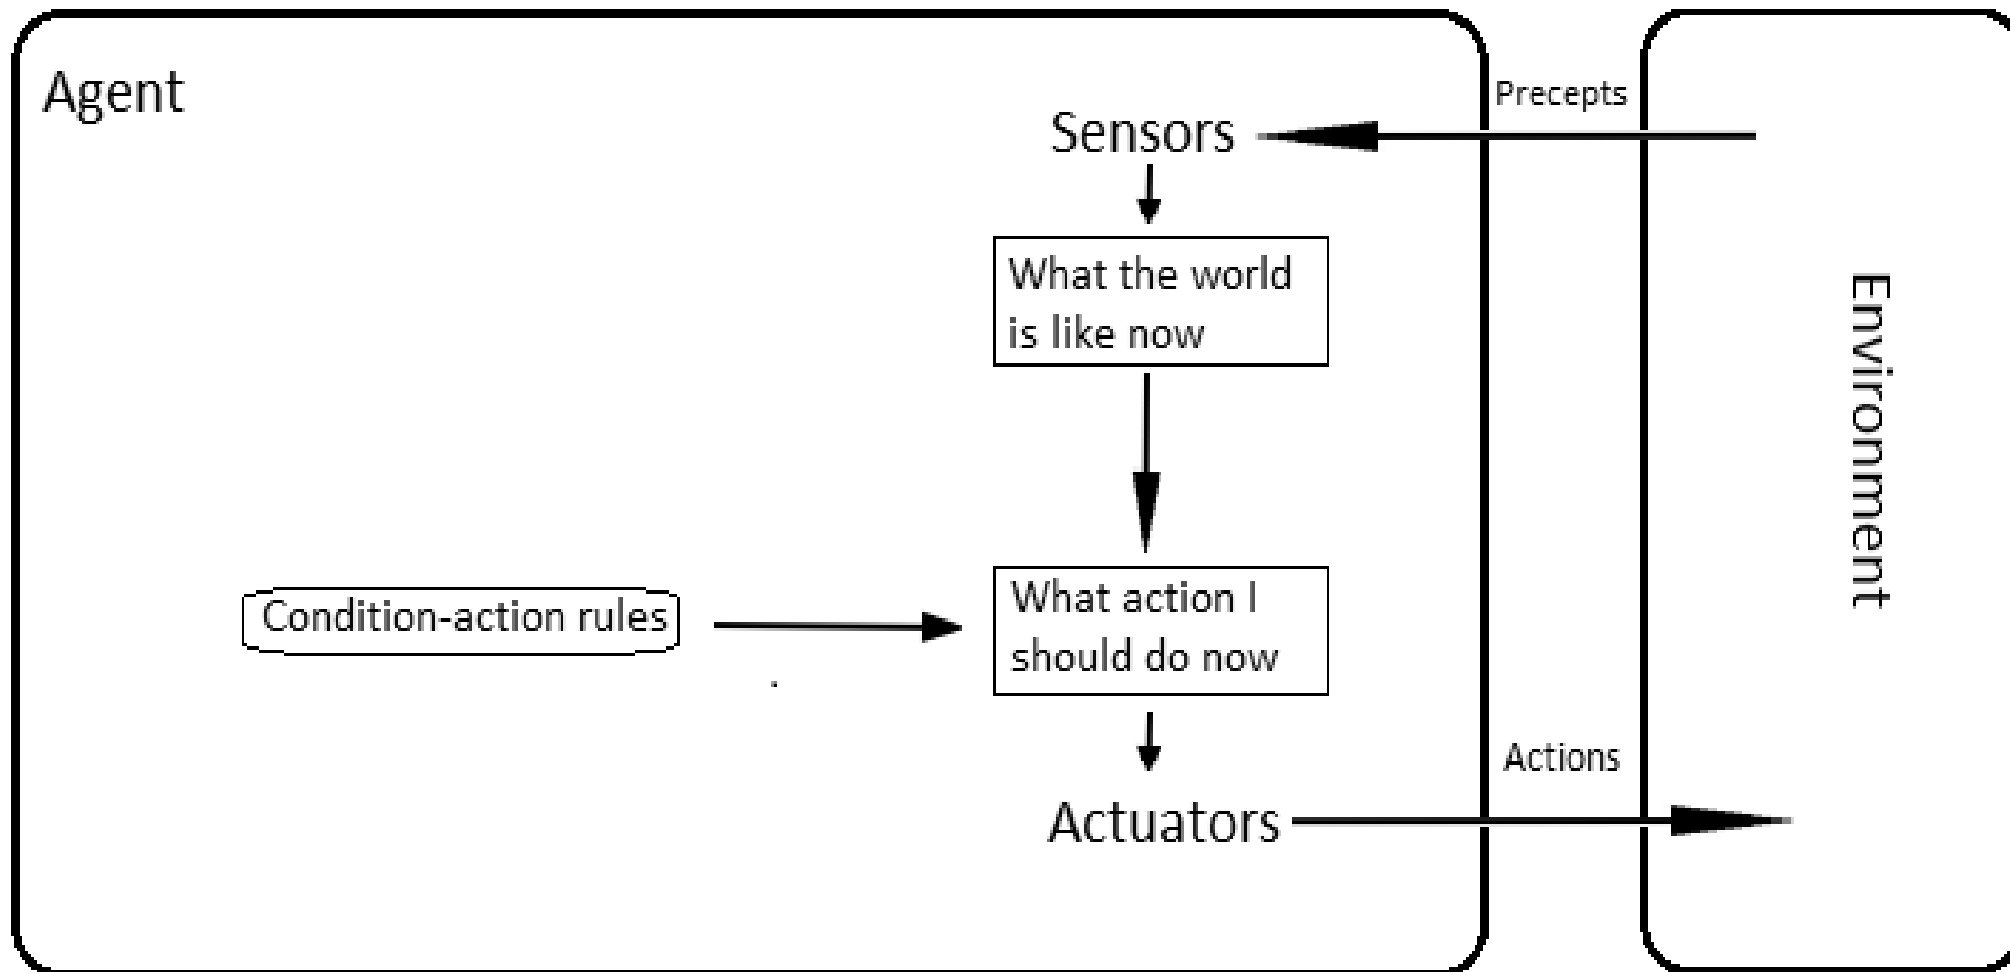
Intelligent Agents

- Agent: Definition Any entity that *perceives* its environment through *sensors* and *acts* upon that environment through *effectors*
- Examples (class discussion): human, robotic, software agents

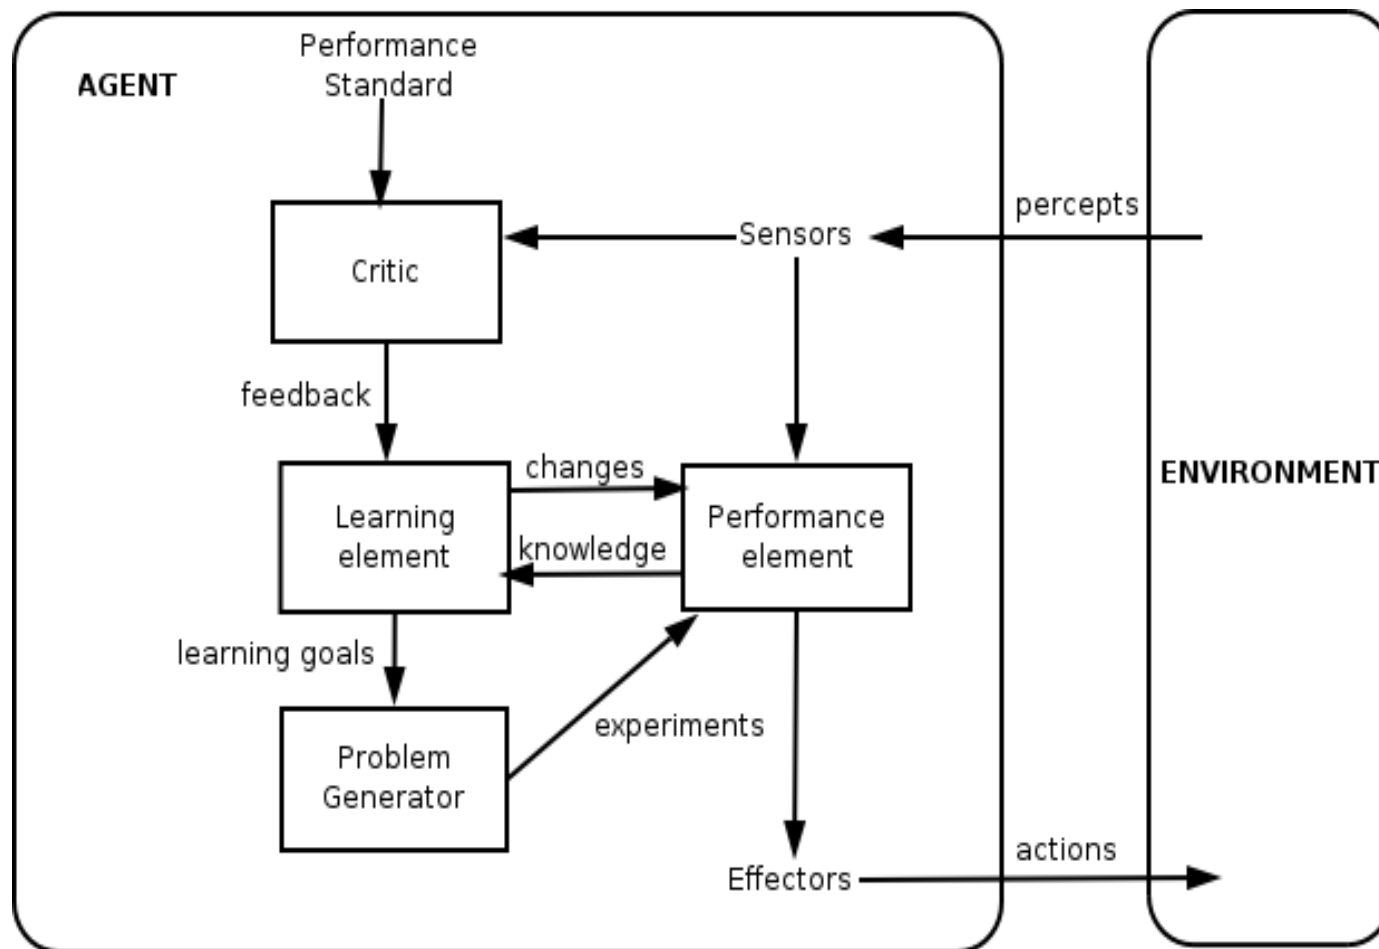
How Agents Should Act

- A rational agent:
 - Does the *right thing*
 - Given what it *believes*
 - From what it *perceives*
- **Maximum Success Measure (Utility)**
 - What is the *right thing*
 - *How* and *When* to evaluate success

Simple Reflex Agents



Learning Agents



Logic Theorist

- Allen Newell, Herbert Simon
- Wrote a program with problem solving skills in 1955
- Used to prove theorems



Methods for Developing AI


- Knowledge Representation
- Search
- Expert Systems & Knowledge Bases
- Planning: classical, universal
- Probabilistic reasoning
- Machine learning: neural networks, evolutionary computing
- Applied AI: agents focus
- Special topics (NLP focus)

Marvin Minsky



Artificial Neural Networks

- Imitation of neurons and their connections
- Each neuron does a task
- Strengthening connections improves the network




Camouflaging Tanks

- 100 photos of tanks behind trees
- 100 photos of just trees
- It worked for all pictures that were used to train the system
- It didn't work for another set of pictures



Why?




Learning to play Mario

Artificial Intelligence Today

Deep Blue





IBM **Watson**[™]



- Google's Alpha Go
 - First AI to beat professional human player (2015)
 - And world champion in 2016
- Go has 10^{170} possible board configurations
- Expanded into AlphaZero

Other Uses

- Microsoft Kinect
- Apple Siri
- Google
- Wolfram Alpha
- Alexa

Almost Everywhere!

In Review

- Alan Turing & the Turing Test
- John Searle & the Chinese Room
- Newell & Simon's Logic Theorist
- Dartmouth Research Project
- Subtopics & Tools in AI
- Marvin Minsky & Neural Networks
- AI Today (briefly)

What are we Missing?

- Philosophical Implications
- Ethical Implications - <https://www.moralmachine.net/>
- Solvability - Is there something a human can do that can't be done by an AI?
- Singularity
- What is "consciousness"?