Artificial Intelligence

CPSC 481

What is Artificial Intelligence?

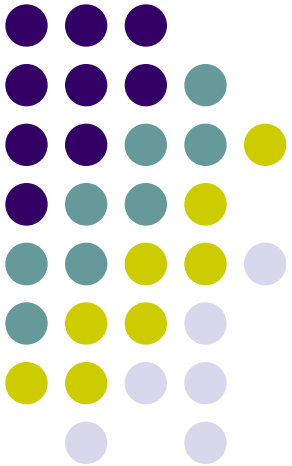

Introduction and Syllabus

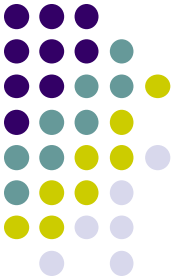

- Waitlist, roster and introduce yourself
- Syllabus and course overview

CPSC 483 Introduction to Machine Learning
CPSC 485 Computational Bioinformatics
CPSC 583 Expert System Design
CPSC 585 Artificial Neural Network

State of the Art Applications

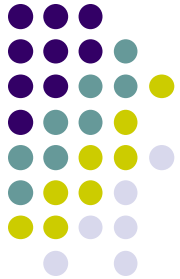

- Board game: AlphaGo and AlphaZero
- Autonomous car: Google Waymo; Tesla autopilot; Audi summon feature; etc.
- Security: Generative Adversarial Network (GAN): Deepfake
- Robot: Sophia; Boston Dynamics; Erica (anthropomorphic)
- Natural language processing-Voice assistant: Google Duplex; Amazon Alexa; Apple Siri; etc.
- Internet of Things (IoT): Amazon cashierless store
- Expert systems: IBM Watson
- Video game: OpenAI-Dota2

History

- 1950 Claude Shannon's "Programming a Computer for Playing Chess" is the first published article on developing a chess-playing computer program.
- 1950 **Alan Turing** publishes "Computing Machinery and Intelligence" in which he proposes "the imitation game" which will later become known as the "Turing Test."
- 1956 the First AI Conference at Dartmouth: A program that simulates the machine; How can a computer be programmed to use a language; Neuron networks; Self-improvement through learning and reasoning
- 1970 The first anthropomorphic robot, the WABOT-1, is built at Waseda University in Japan. It consisted of a limb-control system, a vision system and a conversation system.
- 1972 MYCIN, an early expert system for identifying bacteria causing severe infections and recommending antibiotics, is developed at Stanford University.
- 1997 **Deep Blue** becomes the first computer chess-playing program to beat a reigning world chess champion.
- No major breakthroughs in many years…

Why is AI Difficult?

???

What is Intelligence?

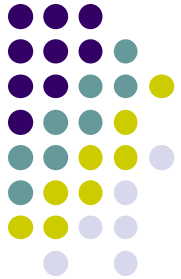

- **Intelligence** is a **general mental capability**.
  - different from being "smart" that is capable of quickly adapting to the environment
- **Some capabilities**
  - learning
  - recognition, understanding
  - reasoning
  - imagination (or abstract thinking)
  - creativity
  - communication
  - planning
  - decision making
  - problem solving

Theories of Intelligence

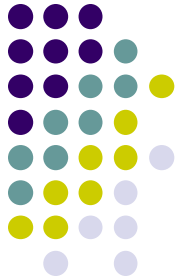

- **Single intelligence**
  - Single general intelligence
    - **IQ** (psychometric testing) to measure the level of intelligence
- **Multiple intelligence**
  - Logical, linguistic, spatial, musical, kinesthetic, interpersonal, intrapersonal, and naturalist intelligences (sensing patterns and making connections to elements in nature) by Howard Gardner
  - Emotional intelligence by Daniel Goleman and others
    - But the Gardners theory has never been tested
    - **EQ** to measure emotional intelligence

Triarchic theory of intelligence

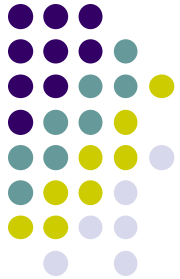

- Formulated by Robert J. Sternberg
- Analytical/componential intelligence
  - take apart problems and being able to see solutions not often seen; “book smart”
- creative/experiential intelligence
  - how well a task is performed with regard to how familiar it is. two parts: novelty and automation. A *novel* situation is one that you have never experienced before. A process that has been *automated* has been performed multiple times and can now be done with little or no extra thought.
- practical/contextual intelligence
  - "deals with the mental activity involved in attaining fit to context" ; "street smarts"

Other Aspects of Intelligence

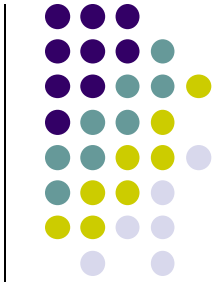

- Animal or plant intelligence
  - Is human intelligence different from animal or plant intelligence?
- Factors affecting intelligence
  - Biological
  - Environmental
- Evolution of intelligence

- Many aspects of intelligence and cognitive process are still unknown
  - Exactly what happens when learning occurs?
  - How is *knowledge represented* in the nerve system (or brain)?
  - What is intuition or skepticism?
  - What is self-awareness?
  - etc.

Artificial Intelligence (AI)

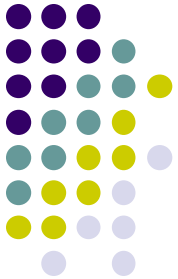

- **AI** influenced by
  - many people in different areas including *philosophers*, *psychologists*, ***mathematicians***, ***computer scientists***, etc.

- **AI** is *Artificial Intelligence* using computer.
  - Artificially made intelligence

- **Other definitions of AI** (as a field of studies)
  - The science and engineering of making intelligent machines (John McCarthy)
  - The study and design of intelligent agents where an intelligent agent is a system that perceives its environment and takes actions that maximize its chances of success.

Intelligent Agents

- An agent is anything that can be viewed as perceiving its environment through sensors and acting upon that environment through actuators
  - Human agent has eyes, ears, and other organs for sensors; hands, legs, mouth, and other body parts for actuators.
  - Robotic agent has cameras and infrared range finders for sensors; various motors for actuators.
- The agent function maps from percept histories to actions: [f: $\mathcal{P}^* \rightarrow \mathcal{A}$]
- The agent program runs on the physical architecture to produce
  - **agent = architecture + program**
- Rational agents use
  - environment,
  - sensors,
  - performance measure,
  - actuators

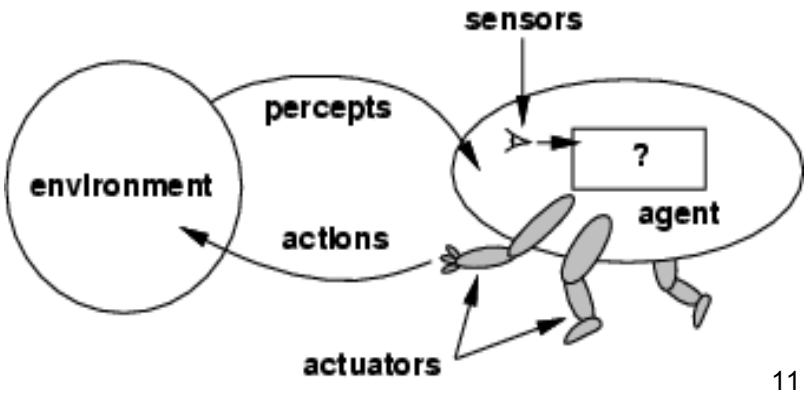

Properties of Task Environments

- **Fully observable vs. partially observable**: – Fully observable if agents sensors detect all aspects of environment relevant to choice of action; Could be partially observable due to noisy, inaccurate or missing sensors, or inability to measure everything that is needed; Often, if other agents are involved, their intentions are not observable, but their actions are;
- **Deterministic vs. stochastic**: Deterministic: If the next state of the environment is completely determined by the current state and the action executed by the agent; the next state has some uncertainty associated with it; Uncertainty could come from randomness, lack of a good environment model, or lack of complete sensor coverage
- **Episodic vs. sequential:** The agent's experience is divided into atomic episodes. The next episode does not depend on the actions taken in the previous episodes. Sequential if current decisions affect future decisions, or rely on previous ones. Examples of episodic are expert advice systems – an episode is a single question and answer.

Properties of Task Environments-cont.

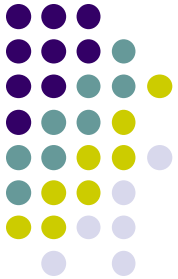

- **Static vs. dynamic**: Dynamic if the environment may change over time. Static if nothing (other than the agent) in the environment changes. Other agents in an environment make it dynamic. Not dynamic if the agent moves from one part of an environment to another, though it has a very similar effect
- **Discrete vs. continuous**: time moves in fixed steps, usually with one measurement per step (and perhaps one action, but could be no action). Playing chess is discrete; Signals constantly coming into sensors, actions continually changing. taxi-driving is continuous.
- **Known vs. unknown**: In a known environment, the outcomes for all actions are given. If the environment is unknown, the agent will have to learn how it works.
- **Single agent vs multiple agents**

Example- playing a single soccer game

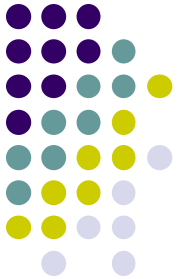

- Partially observable – An agent cannot detect all the things on soccer field that can affect its action, for e.g. it cannot determine what other players are thinking.
- Stochastic – For a given current state and action executed by agent, the next state or outcome cannot be exactly determined, for e.g., if agent kicks the ball in a particular direction, then the ball may or may not be stopped by other players, or the soccer field can change in many different ways depending on how players move.
- Sequential – The past history of actions in the game can affect the next action in the game.
- Dynamic – The environment can change while the agent is making decision, for e.g., soccer field (environment) changes when a player moves.
- Continuous – Location of the ball or players is continuous. The speed or the direction (angle) at which the agent hits the ball is continuous.
- Unknown
- Multi-agent – There are many agents involved in soccer game.

Intelligent Agents

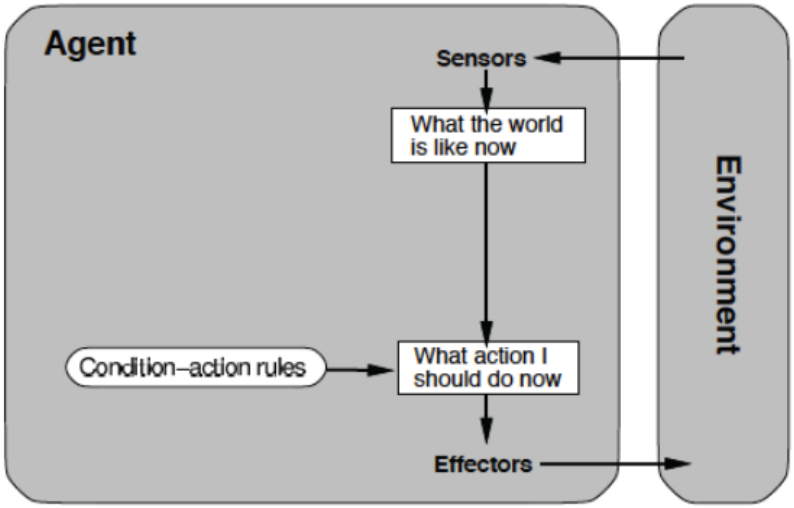

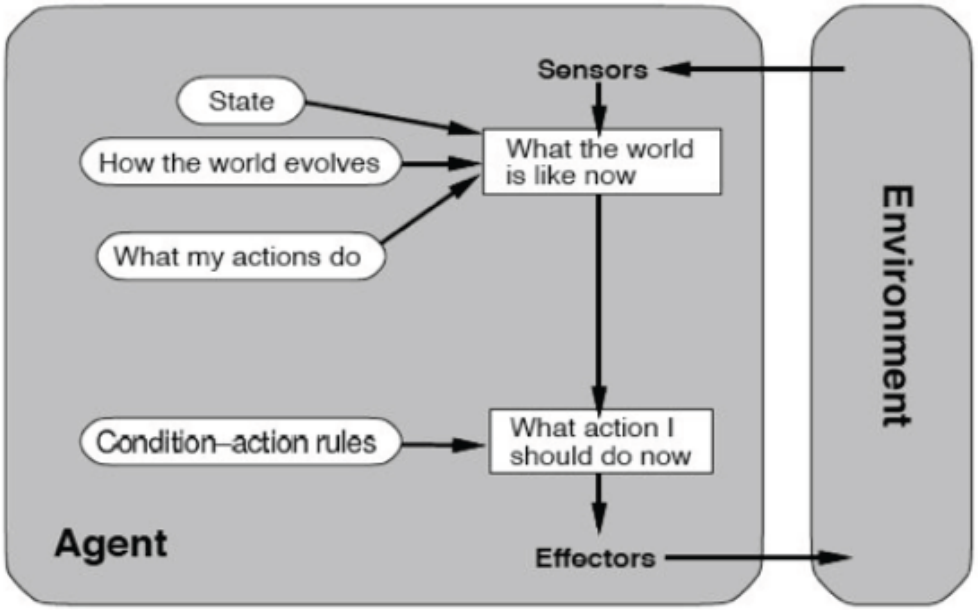

Simple reflex agent: only on the basis of the current percept, ignoring the rest of the percept history. only succeeds when the environment is fully observable. Infinite loops in partially observable environments.

Simple reflex agent with internal state (model-based agent): handle partially observable environments. Its current state is stored inside the agent maintaining some kind of structure which describes the part of the world which cannot be seen.

Intelligent Agents with Goals

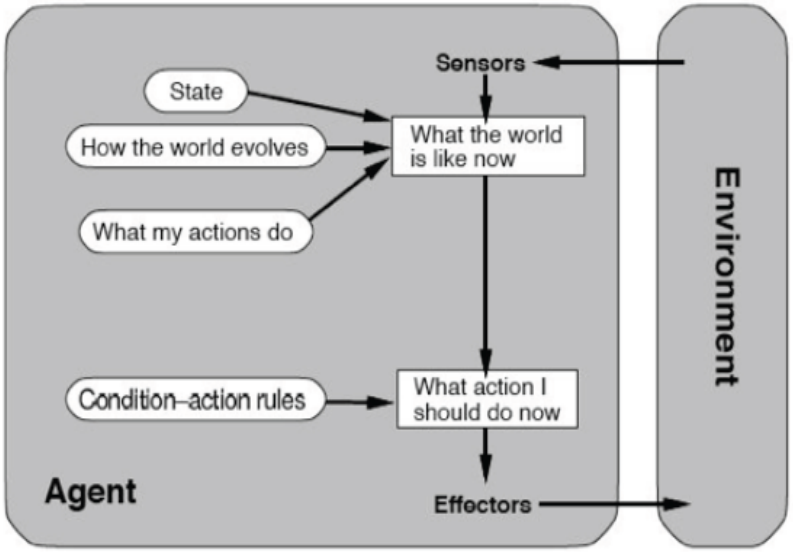

Simple reflex agent with internal state **(model-based agent)**

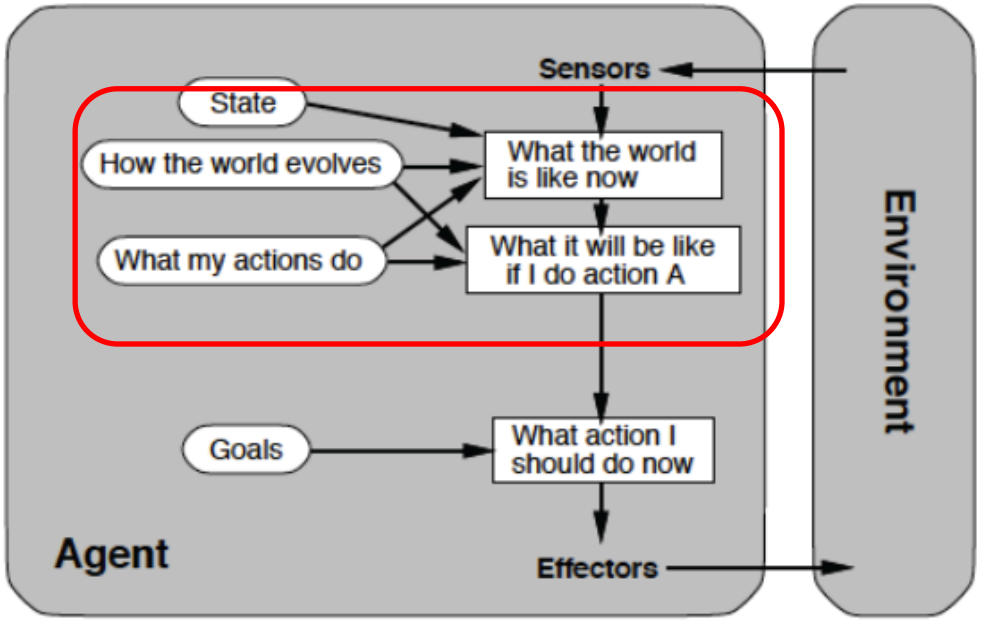

Goal-based agent that has reasoning/inference and **goals:** This allows the agent a way to choose among multiple possibilities, selecting the one which reaches a goal state.

Intelligent Agents with Utility

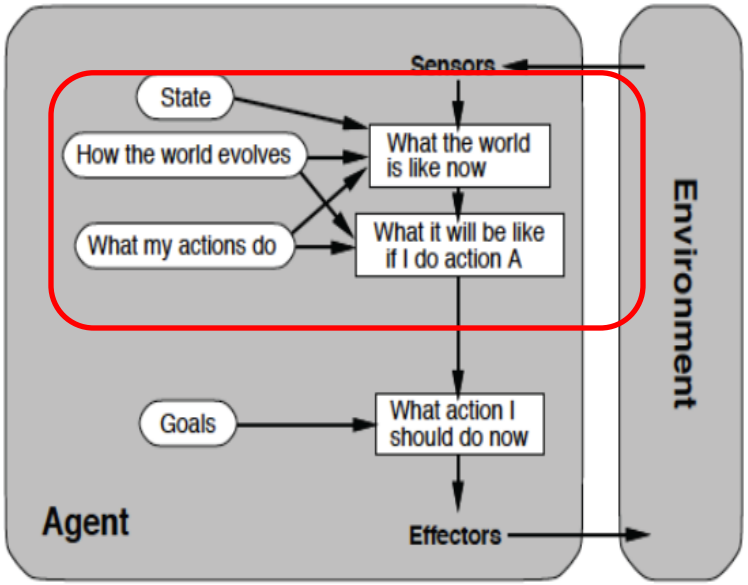

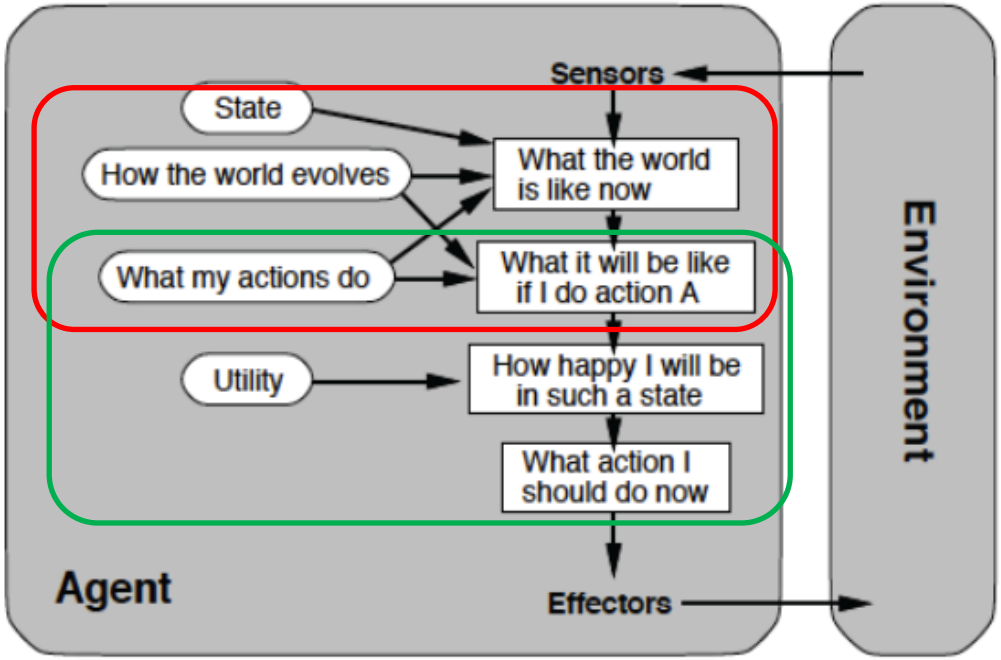

Goal-based agent that has Reasoning/inference and **goals**

Utility-based agent that consists of: reasoning/inference and rational decision making

Major Areas of Interests in the First AI Conference (at Dartmouth, 1956)

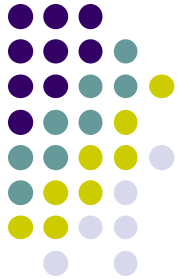

- Automatic computers
  - A program that simulates the machine
- How can a computer be programmed to use a language
  - The way that human uses and manipulated words
- Neuron networks
- Theory of the size of a calculation
- Self-improvement through learning and reasoning
- Abstractions
- Randomness and creativity
  - Difference between randomness and creativity

General Areas of AI

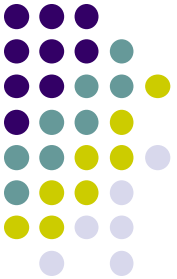

- AI research is specialized and deeply divided into subfields
  - Perception and the ability to move and manipulate objects
  - Knowledge representation
  - Natural language processing and understanding
  - Machine learning
  - Reasoning: Knowledge-based or Expert systems
  - Planning
  - Problem solving
  - Complex decision making
  - Abstract thoughts or creativity
- Scope of AI
  - *What system is considered intelligent?*

Directions of Artificial Intelligence

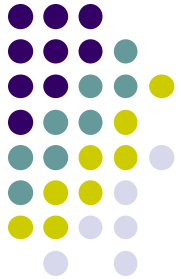

- **Weak AI** – AI with specific intelligent abilities
- **Strong AI** – AI with general intelligence

- **Directions of AI research**
  - From "Thinking humanly and acting humanly"
  - To "Thinking rationally and acting rationally"
    - (through ideal decision making)

Turing Test

(as a Way of Testing Acting Humanly)

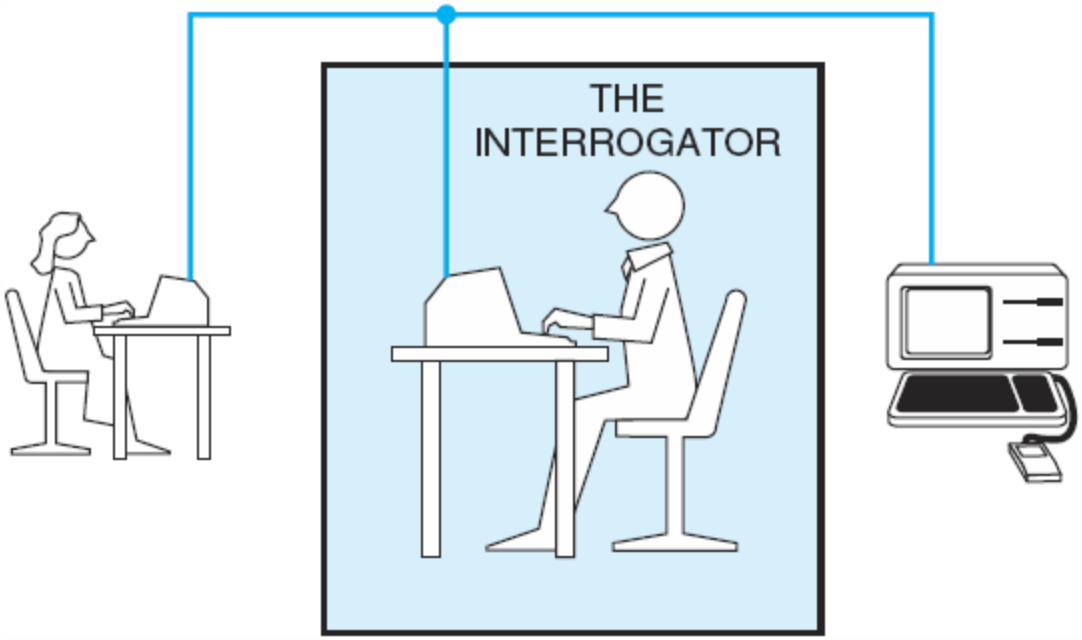

Turing test by Alan Turing in the theory of computability

Asking a question “Can a machine to be made to think?”

The Turing Test measures the performance of an intelligent machine against that of a human being through the imitation game.
Trying to determine which player – A or B – is a computer and which is a human

Can a system that passed a Turing Test be considered as intelligent as human?

Chatbots: ELIZA (1966), social media messengers, etc.

Importance and Criticisms for Turing Test

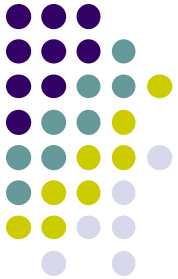

- **Importance of Turing Test**
  - Attempts to give an objective notion of intelligence by eliminating any bias in favor of living organisms.
- **Criticisms for Turing Test**
  - Focuses on purely symbolic and problem-solving skills, not on perceptual skill.
  - Computers can only do as they are told and consequently cannot perform original actions (Ada Lovelace).
  - A distraction to the important tasks such as developing general theories to explain the mechanisms of intelligence in human and machines to solve problems.
  - Only based on simulating human intelligence. Machine can do a lot better than human for many tasks.

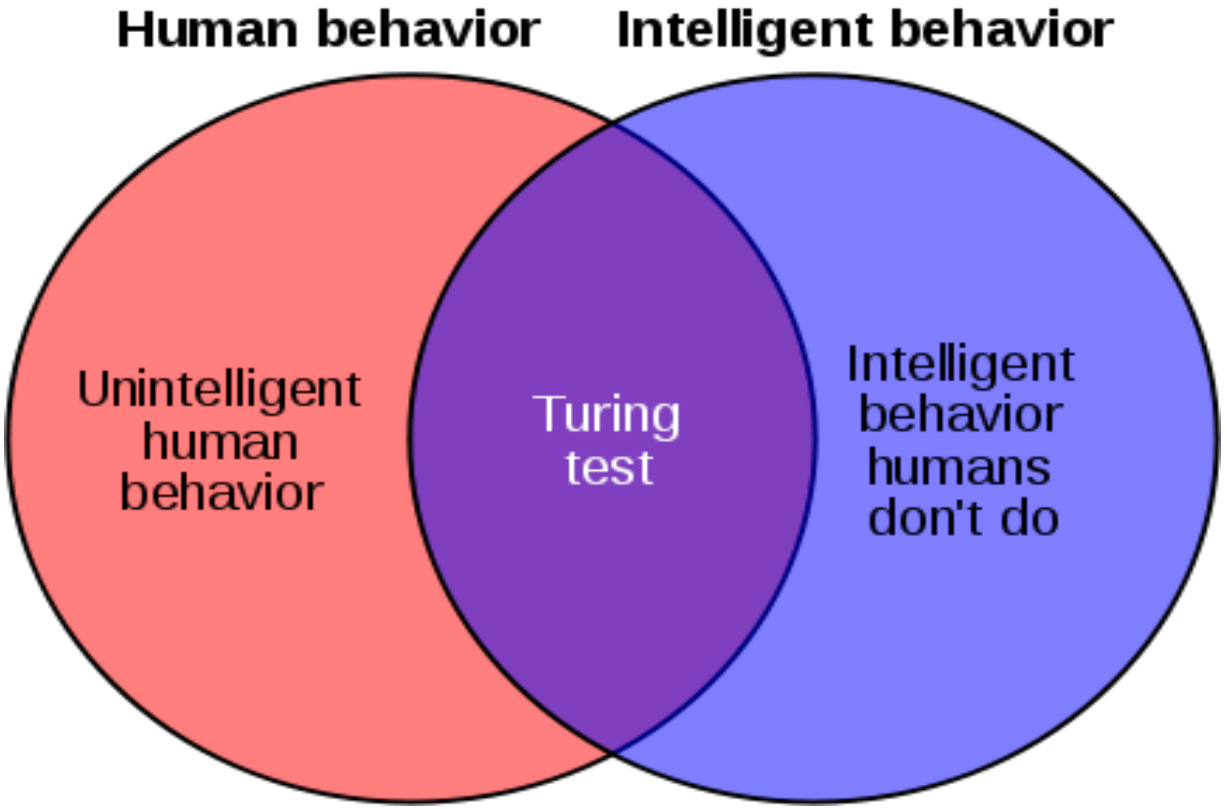
Human behavior
Intelligent behavior
Unintelligent human behavior
Turing test
Intelligent behavior humans don't do

Evolution of AI

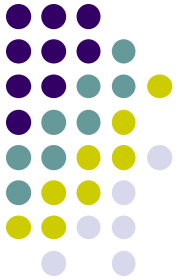

- **AI** has been the subject of optimism, but has also suffered setbacks and, **today and future**, has become an *essential part of the technology in industry*.
- **Early stages**
  - Focused on concepts, theories, and simulating human intelligence such as formal logics, theorem proving, graph theory, search, knowledge representation, stochastic analysis, other related mathematics.
- **Modern AI**
  - Focus more on practical applications and problem solving – demand from real-world business applications.
    - Machine learning, data mining, search engines, intelligent systems, robots, big data analytics, etc.
  - Collective intelligence

Future of AI

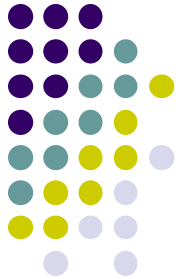

- AI started with the goal of strong AI but it has been weak AI.
  - *AI (robot or agent with AI) still weak at recognition, feeling, emotion*
  - **Strong AI** is still a long-term goal of AI research.
  - AI offers a unique and powerful tool for expanding the capabilities of computer science.
- **Questions to think about**
  - Can a machine equipped with strong AI be more intelligent than human?
  - What ethical issues may arise when we have the true intelligent system with strong AI?
- How can Human and AI coexist?
  - Human player utilizing computer chess could beat Deep Blue.
- Future job: How well can you utilize AI?
  - We are more influenced by pessimist's forecasting.
  - **What has happened**: History tells that different jobs will be created.

Primary Focus in this Course

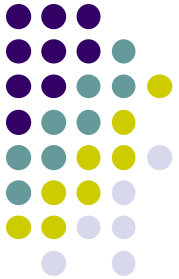

- While we are trying to *understand intelligence and implement intelligence* such as **problem solving**, **reasoning**, **learning**, **creativity**, natural language understanding, decision making, and planning,

- We also want to learn practical AI as complex and intelligent problem solving strategies.

Study Guide

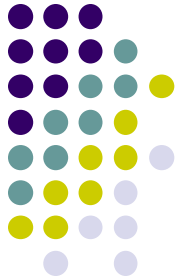

- **Levels of knowledge capabilities**
  - Being able to understand and explain (can tell)
  - Being able to apply (can do)
  - Being able to teach (can teach)
  - Being able to criticize (can criticize)
  - Being able to improve and innovate (can innovate)

- Think about the review questions and subjects related to the questions and make sure you ***understand the principle ideas***.
- Once you understand them, think about how you can **apply** them to solve complex problems and build an intelligent system.

Review Questions -1

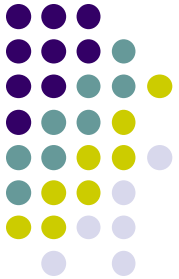

- What is intelligence?
- What are the key capabilities to being intelligent? Try to define the meaning of each capability and think about potential real-world applications.
- What is the main difference between single intelligence and multiple intelligence?
- Does only human has intelligence?
- What are some unknown aspects of intelligence?
- Can we develop a software system that has intelligence?
- What is Artificial Intelligence (AI)?
- What is an intelligent agent?
- What are the different types of environment for agents?
- Define simple reflex, goal-based, and utility-based agent.
- What are the areas of AI?
- How can we tell if a software system is intelligent?
- How can AI be used for solving complex real-world problems? What are the applications of AI?

Review Questions -2

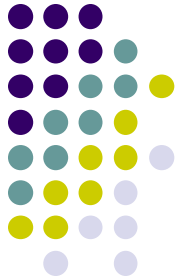

- Is it possible to develop a system that is as intelligent as human?
- What is Turing test? What is the purpose of Turing test?
- Is it possible to develop a system that can pass Turing test?
- If a machine passes the Turing test, can the machine solve most AI problems?
- Is it possible to develop a system that is more intelligent than human?
- What is the main difference between weak AI and strong AI?
- What do you think will be the future of AI?
- What may be potential problems we may encounter with advanced AI?

Most Important Points to Remember

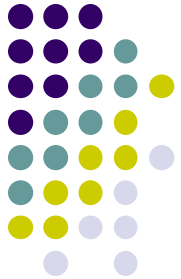

- Can you explain about intelligence?
- What are the major capabilities for intelligence?
- Can you explain about artificial intelligence (AI)?
- How can we tell if a system is intelligent?
- What is an intelligent agent?
- What are the areas of AI?
- What are the applications of AI in real-world?
- What is the ultimate goal for AI?

References

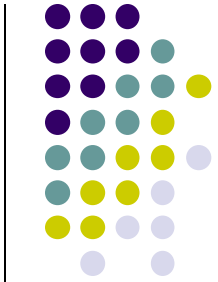

- George Fluger, Artificial Intelligence: Structures and Strategies for Complex Problem Solving, 6th edition, **Chapter 1**, Addison Wesley, 2009.
- Russel and Norvig, Artificial Intelligence: A Modern Approach, 3rd edition, Prentice Hall, 2010.