

Artificial Intelligence

- More than any other subfield of computer science, *artificial intelligence*, or AI for short, has probably caught the imagination of the general public the most
- Amidst science fiction visions of droids, robots, and evil computers taking over the world, the real-world subfield of AI includes a variety of concepts and endeavors, each of which hopes to enable machines to exhibit behavior that we would classify as “intelligent”
- AI is also highly multidisciplinary, integrating psychology, neurology, linguistics, and other fields

Central Concept: Agents

- AI focuses on the notion of an intelligent *agent* — a “device” (not necessarily physical) that responds to stimuli in a manner that we interpret as “intelligent”
- AI seeks to develop agents that respond to their stimuli in a sensible and rational manner
 - ◇ *Perception* enables agents to “understand” or process incoming data effectively
 - ◇ *Reasoning* handles the formation of appropriate responses to the perceived stimuli
- Another aspect of AI explores *learning* — the ability for an agent to *improve upon* its behavior over time, either through *procedural* or *declarative knowledge*

The Turing Test

- The notion of “intelligent” agents begs the question — what counts as “intelligence?”
- An early (and still somewhat popular) benchmark was proposed by (who else) Alan Turing in 1950
- The idea behind the *Turing test* is to have an *interrogator* converse with another entity purely by typing and reading (not unlike today’s instant messaging)
- If the interrogator could not tell whether the entity was a human or machine, then the behavior exhibited can be viewed as “intelligent”

Strong vs. Weak AI

- The Turing test is based on the idea that intelligence is based on external *behavior*; however, some argue that intelligence lies in *consciousness* or *sentience*
- The *weak AI* camp contends that humans are fundamentally different from machines, and that machines can only *act* intelligently, not *be* intelligent
- The *strong AI* camp views that machines *can* ultimately achieve consciousness — it’s a matter of making the right parts and putting them together in the right way
- Where would *you* stand?

Perception

- Regardless of the final result (strong AI vs. weak AI), certain problems remain to be solved no matter what
- *Perception* presents some key challenges: how can we make an agent “sense” its surroundings?
- Current technology has produced a wide variety of *sensors* that can quantify and digitize a variety of real-world phenomena:
 - ◆ *Light sensors* can measure and digitize brightness and color
 - ◆ *Sound sensors* can measure and digitize sound pitch and volume; pairs of sound sensors at ultrasonic frequencies use *triangulation* to measure distances to physical objects
 - ◆ *Touch and motion sensors* can measure and digitize physical pressure, speed, and movement

- While machines can have a range of artificial “sensory organs,” they continue to lack the ability to *interpret* or *understand* these digitized sights and sounds
- *Image understanding* seeks to develop algorithms that facilitate machine interpretation of raw digital images into meaningful objects, shapes, patterns, and symbols
 - ◆ Image understanding can be viewed as two steps: *image processing* (manipulating and transforming pictures) and *image analysis* (deriving meaning from the resulting content)
- *Natural language processing (NLP)* is concerned with algorithms that can derive meaning from the languages that we use — English, Spanish, Latin, etc.
 - ◆ NLP also has a number of accepted phases: *syntactic analysis* identifies the roles taken by words and phrases; *semantic analysis* attaches meanings to the roles; and *contextual analysis* integrates associated information (surrounding words, real-world knowledge)
 - ◆ NLP assumes that the words are already represented digitally (i.e., text) — to understand the spoken word, we need *speech recognition* to translate digitized sound into text first

Reasoning

- Once an agent's environment has been adequately perceived, the agent must then do some *reasoning* in order to make decisions, solve problems, or exhibit intelligent behavior
- Like perception, this area of AI also encompasses a variety of distinct approaches
- *Production systems* represent problem solving as a set of *states*, one of which is the *goal state*, over which a *control system* applies rules or *productions* that move from one state to another, eventually reaching the goal

- *Inference engines* constitute a form of production system, with sets of assertions making up a state, and *inference rules* serving as productions for deriving new assertions based on the previous ones
 - ◆ Inference engines form the core of *expert systems*, which seek to simulate human reasoning for a specific domain or problem area
- Drawing out a production system's transitions from state to state yields a *directed graph* of all possible states; starting from an initial state outward forms a *search tree* of possibilities as we try to reach the goal
- Still, many problems remain too large (chess!) for a system to exhaustively search every single possibility; to narrow our choices, systems often apply *heuristics* that quantifies how "good" a decision might be

Knowledge Bases

- An issue that underlies both perception and reasoning is the concept of *knowledge* — what is it, and how can it be effectively represented within a machine?
- We can represent numbers, words, pictures, and sounds as bits...but how does one represent something like “Being sick reduces productivity” or “High reward sometimes motivates high risk”
- Beyond representation, can we make machines identify *implied* knowledge, or even *derive* or *extract* knowledge from pre-existing information?

[Machine] Learning

- Thoughts on representing and manipulating knowledge naturally leads to the idea that machines might be able to *acquire* new knowledge — a.k.a. *machine learning*
- A wide variety of approaches have been tried:
 - ◆ With *imitation*, a machine “watches” and records how a person accomplishes a task
 - ◆ *Supervised training* takes human choices that are made for some *training set* and uses these choices to influence the machine’s later behavior and decisions
 - ◆ In *learning by reinforcement*, machines are given a method to evaluate success or failure; the machine undergoes a trial-and-error phase and tracks successful attempts
 - ◆ *Genetic algorithms* try to develop solutions by “evolving” them from a pool of possibilities — candidate algorithms are repeatedly pruned then “re-mixed” into new “generations”

Artificial Neural Networks

- One approach to AI focuses on the nature of the machine itself: since we are trying to replicate the abilities of the human brain, then why not build a machine that resembles it?
- *Artificial neural networks* (ANNs) try to mimic the known structure of the human brain: interconnected cells that hold electrical charges and are influenced by the charges of the cells that surround them
- ANNs aren't explicitly programmed but *trained*: test scenarios adjust how cells affect each other

Associative Memory

- The human brain is capable of *association* — moving quickly from one idea to another, such as a song to fond memories, or a scent to a particular individual
- This observation leads to *associative memory*, which is another approach to mimicking the brain
- Associative memory can be implemented using ANNs; the approach is to define “stable” and “unstable” cell values, and to let the network transition from one stable state to another in response to some input

◆ The idea here is to make certain input lead the ANN from one state to another, similar to how catching a scent may remind us of someone we know

Robotics

- Recall that many agents don't have a physical presence — they can be pure software, such as a chess-playing program or an expert system
- The area of *robotics* explores the idea of *physical* and possibly *autonomous* agents in the real world
- Robots (from the Russian *robota*, for “work”) respond physically to physical stimuli (motion, sound, lights)
- All prior AI notions apply to robots, with the addition of mechanical subsystems for real-world interaction, which may or may not mimic the human body