Chapter 12: Artificial Intelligence and Modeling the Human State

Are computers smart enough to replace people?

12-1

What is Artificial Intelligence?

- What is your concept of AI?
- One definition
  - AI is the study of how to make computers do things that people (generally) do better.
  - Mundane
    - Vision, Speech
    - Natural Language Processing, Generation, Understanding
    - Reasoning
    - Motion
  - Formal
    - Board Game-Playing, chess, checkers
    - Logic
    - Calculus
  - Expert
    - Design, engineering, graphics
    - Financial Analysis

12-2

What is Intelligence: Artificial or Not?

- Does “looking intelligent” mean that intelligence is present?
  - **Maillardet’s Automaton** (Henri Malliardet, 1805):
    - Object having human form seemed to mimic the intelligence of the human.
    - Drawing machine.
      - Disguised as a young boy.
      - Containing levers, ratchets, cams and other mechanical devices.
      - Could draw several complex images.
    - Because it had human form and could draw complex images, a certain feeling of intelligence was ascribed to the machine.

12-3

What is Intelligence: Artificial or Not?

- Sailing vessel drawn by Maillardet’s Automaton.

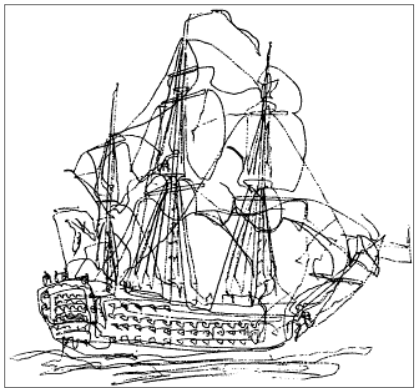

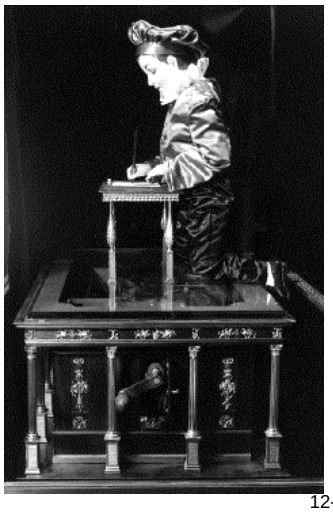

12-4

What is Intelligence: Artificial or Not?

- Alan Turing (1912 - 1954)
  - Proposed a test - **Turing's Imitation Game**
    - Tests the intelligence of the computer.
  - Phase 1:
    - Man and woman separated from an interrogator.
    - The interrogator types in a question to either party.
    - By observing responses, the interrogator's goal was to identify which was the man and which was the woman.

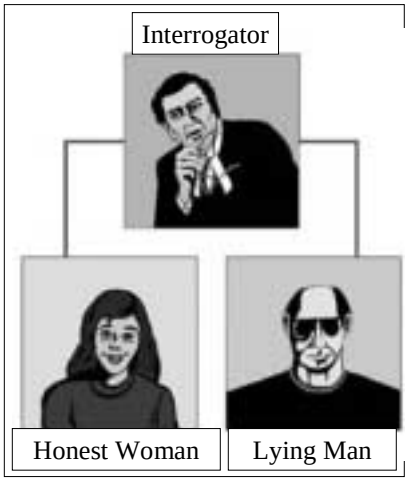

12-5

What is Intelligence: Artificial or Not?

- Phase 2 of the Turing's test:
  - The man was replaced by the computer.
  - If the computer could fool the interrogator as often as the person did, it could be said that the computer had displayed intelligence.

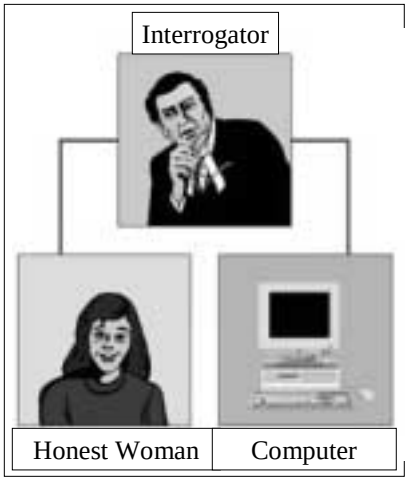

12-6

Philosophical Foundations

- Do you think the Turing Test is a good measure of intelligence?
- Some counter-arguments
  - Searle's Chinese Room
    - Room with a slot, human with huge rule book on how to translate Chinese to English
    - If someone drops a Chinese letter in the slot and the human translates it to English, does the human understand Chinese?
  - Measures imitation, not intelligence
    - Eliza program designed to fool, not "think"

12-7

Natural Language Communication

- The semantics (meaning of the words) is the goal in communication with the computer.
  - ELIZA (1965) by Joseph Weizenbaum of MIT
    - Simulation of a therapist mirroring a patient's conversation in a positive manner.
    - The program fooled people into believing it was a therapist.
  - Uses rules like:
    - "I like _____"
    - Eliza could respond, "Really? I hate _____, it's terrible."
    - Seems human-like, but it is really just matching patterns
    - Lots of these together makes the system seem more intelligent
  - Loebner contest
    - $100,000 to winner that can pass the Turing Test

12-8

Natural Language Communication

- Part of an ELIZA Conversation:

All men are alike.
IN WHAT WAY
They're always bugging us about something or other.
CAN YOU THINK OF A SPECIFIC EXAMPLE
Well, my boyfriend made me come here.
YOUR BOYFRIEND MADE YOU COME HERE
He says I'm depressed much of the time.
I AM SORRY TO HEAR YOU ARE DEPRESSED
It's true, I am unhappy.
DO YOU THINK COMING HERE WILL HELP YOU NOT TO BE UNHAPPY
I need some help, that much seems certain.
WHAT WOULD IT MEAN TO YOU IF YOU GOT SOME HELP
Perhaps I could learn to get along with my mother.
TELL ME MORE ABOUT YOUR FAMILY

12-9

Natural Language Communication

- Semantic Translation Problems (Problems with language translation).
  - A classic example, called the Bar-Hillel paradox, illustrates a difficult semantic problem:

    The pen is in the box.
    The box is in the pen.

    – Both sentences have identical syntax structures.
    – Interpretations:
      - First statement: A writing instrument is in the box.
      - Second statement: A box is in the playpen.
  - Convinced Bar-Hillel that computer translation of languages was impossible.
- Solution : Need way to track context, represent lots of domain knowledge!

12-10

Natural Language Communication

- Early attempts at language translation:
  - An early attempt to translate an English expression to Russian and back again to English:
    - Typed in English (sentence to be translated...):
      - The spirit is willing, but the flesh is weak.
    - Translated by the program into Russian and back into English:
      - The vodka is strong, but the meat is rotten.
- Translation programs have come a long way.
  - WWW translation programs
    - Accuracy and interpretation still very crude.
  - Look for multiple definitions of words, try to find a way to make them match up grammatically and somewhat semantically
  - Lacks the necessary domain knowledge to ensure the translation makes sense pragmatically

12-11

Natural Language Translation

- Web-based Language Translation
  - Babel Fish (Free service on Alta Vista)
    - Text is cut and then pasted into a translation box.
    - "Test translation" from English to Italian and back:
      - The spirit is willing, but the flesh is weak.
      - The spirit is arranged, but the meat is weak person.
  - FreeTranslation.com
    - Allows you to enter a URL and then translates it.
    - Also does text entry for direct translation to and from English.
    - "Test translation" from English to German and back:
      - The spirit is willing, but the flesh is weak.
      - The intellect is ready, but the meat is weak.

12-12

More Philosophical Issues

- Physical Symbol Hypothesis
  - Newell & Simon, 1976
  - The thinking mind consists of the manipulation of symbols. That is, a physical symbol system has the necessary and sufficient means for general intelligent action.
- If true, then a computer has the necessary means to implement general intelligent action
- Counter-arguments
  - Lack of consciousness
  - Lack of self-awareness
  - Chalmers Mind Experiment

12-13

Branches of AI

- Three major branches of AI
- Strong AI
  - The study and design of machines that simulate the human mind to perform intelligent tasks. Borrows many ideas from psychology, neuroscience, etc. The goal is to perform tasks the way a human might do them, but implement it on the computer.
- Weak AI
  - The study and design of machines that perform intelligent tasks. Not concerned with how tasks are performed, mostly concerned with performance and logic. E.g., to make a flying machine, use logic and physics, don't mimic a bird.
- Emergent AI
  - The study and design of machines that simulate simple creatures, and attempt to evolve and have higher level emergent behavior

12-14

Modeling Human Intelligence

- Modeling human intelligence systems:
  - One way to study complex systems is to build a working model of the system, and observe it in action.
  - Two (of several) approaches to model some of the thinking patterns of the human brain:
    - **Semantic networks**
    - **Rule-based systems** or **Expert systems**

12-15

Modeling Human Intelligence

- **Semantic networks** are designed after the psychological model of the human associative memory.

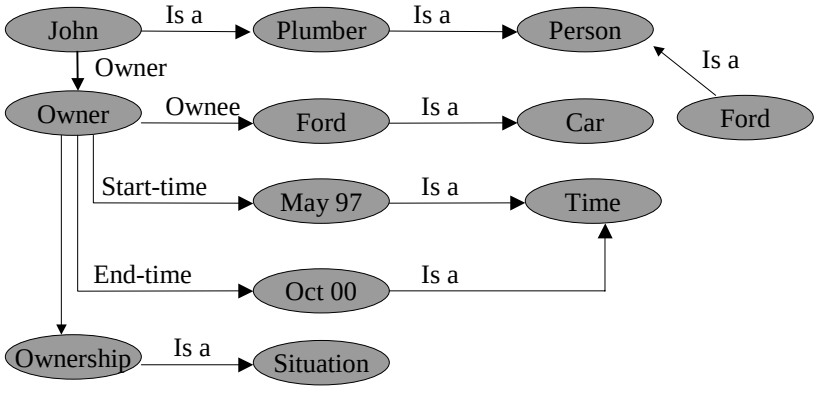

12-16

Modeling Human Intelligence

- **Rule-based or Expert systems** - Knowledge bases consisting of hundreds or thousands of rules of the form:

  IF (condition) THEN (action).

  - Use rules to store knowledge ("rule-based").
  - The rules are usually gathered from experts in the field being represented ("expert system").
    - Most widely used knowledge model in the commercial world.

    - IF (it is raining AND you must go outside)
    - THEN (put on your raincoat)

  - Rules can fire off a chain of other rules

    - IF (raincoat is on)
    - THEN (will not get wet)

12-17

Expert Systems

- **Expert systems** were commercially the most successful domain in Artificial Intelligence.
  - Somewhat out of favor today
  - These programs mimic the experts in whatever field.

Auto mechanic	Telephone networking
Cardiologist	Delivery routing
Organic compounds	Professional auditor
Mineral prospecting	Manufacturing
Infectious diseases	Pulmonary function
Diagnostic internal medicine	Weather forecasting
VAX computer configuration	Battlefield tactician
Engineering structural analysis	Space-station life support
Audiologist	Civil law

12-18

Expert Systems

- Expert systems are also called Rule-based systems.
  - Expert's expertise is built into the program through a collection of rules.
  - The desired program functions at the same level as the human expert.
  - The rules are typically of the form:
    - If (some condition) then (some action)
    - Example: If (gas near empty AND going on long trip) then (stop at gas station AND fill the gas tank AND check the oil).
  - EXCON: An expert system used by Digital Equipment Corp. to help configure the old VAX family of minicomputers.

12-19

Expert Systems

- Two major parts of an expert system:
  - The **knowledge base**: The collection of rules that make up the expert system.
  - The **inference engine**: A program that uses the rules by making several passes over them.
    - On each pass, the inference engine looks for all rules whose condition is satisfied (*if* part).
    - It then takes the action (*then* part) and makes another pass over all the rules looking for matching condition.
    - This goes on until no rules' conditions are matched.
    - The results are all those action parts left.

12-20

Expert Systems

- Inference engines can pass through the rules in different directions:
  - Forward chaining: Going from a rule's condition to a rule's action and using the action as a new condition.
  - Backward chaining: Goes in the other direction.
    - Example: Medical doctors use both.
      - Forward chaining: Going to the doctor with symptoms (stomach pain). The doctor will come up with a diagnosis (ulcer).
      - Backward chaining: The doctor asks if patient has been eating green apples knowing green apples cause stomach aches.

12-21

Modeling Human Intelligence

- For any of these models of the human knowledge system to work, it must be able to make use of this human knowledge in three different ways:
  - **Acquisition** - Must be some way of putting information or knowledge into the system.
  - **Retrieval** - Must be able to find knowledge when it is wanted or needed.
  - **Reasoning** - Must be able to use that knowledge through "thinking" or reasoning.

12-22

Modeling Human Intelligence

- Knowledge Acquisition:
  - A fact is the simplest type of knowledge that can be acquired.
    - Bees sting.
  - Ideas, concepts, and relationships are more difficult for humans and machines.
    - Provoking bees causes them to sting.
    - What ***is*** a chair?
    - Quickly balloons into a huge knowledge representation problem, too much to represent in a computer

12-23

Modeling Human Intelligence

- Knowledge Retrieval by Searching
  - After knowledge has been acquired and stored in one's memory, it can be retrieved and used to solve problems.
  - **Brute-force search** - Looks at every possible solution before choosing among them.
    - Hexapawn game example: The program searches through all the possible moves and then selects the best.
    - The space of possible moves is called the **state space**

12-24

Modeling Human Intelligence

- Hexapawn Game Tree
- Shows different moves ("mirror images" are not shown.)

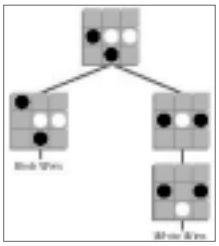

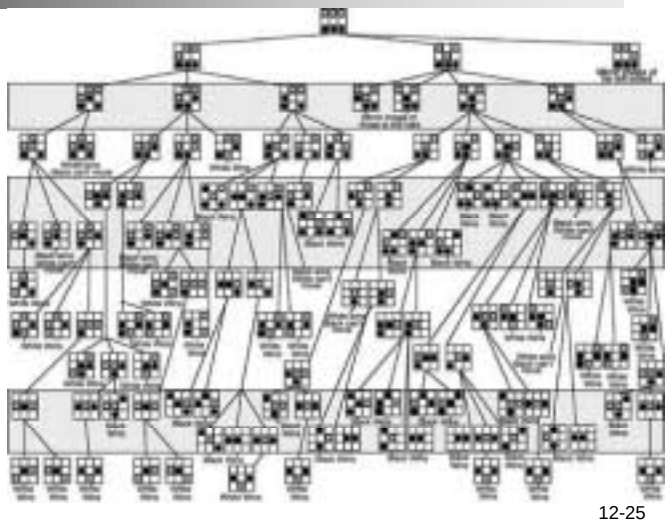

12-25

Modeling Human Intelligence

- **Heuristic search** - **Rules of thumb**, which are used to evaluate a particular state when searching for a solution to a problem. (Not guaranteed to lead to a solution.)
  - Chess game tree would have 10^{120} possible moves.
    - Uses rules of thumb to reduce the number of possible plays.
      - Example: Examine a few plays ahead instead of all the ways to the end of the game.
      - Need some heuristic to evaluate the "goodness" of each state and pick the best one
    - Deep Blue (1996) by IBM - Garry Kasparov, world-champion chess player, won over Deep Blue 4 points to 2.
    - Deep Blue (1997) by IBM - Garry Kasparov conceded victory to Deep Blue, 3.5 points to 2.5.

12-26

Modeling Human Intelligence

- Reasoning with knowledge
  - Humans: Reasoning is what we do when we solve problems.
  - In Artificial Intelligence: Two types of reasoning are commonly used.
    - Shallow reasoning: Based on heuristics or rule-based knowledge.
      - Computers, for the most part, do shallow reasoning.
    - Deep reasoning: Deals with models of the problem obtained from analyzing the structure and function of component parts of the problem.
      - Humans commonly apply deep reasoning.
      - E.g., find an analogy between computational processes and biological processes to better the understanding

12-27

Modeling Human Intelligence

- How can the knowledge base be built up so that there is sufficient knowledge to reason with?
  - Learning systems: Intelligent computer programs that are capable of learning.
  - Types of learning that are used to write intelligent programs:
    - **Rote learning** - Memorization of facts.
    - **Learning by instruction** - Similar to student/teacher relationship found in classrooms.
    - **Learning by deduction** - Drawing conclusions from certain premises (This is a cat. All cats are animals. Therefore, this is an animal.)
    - **Learning by induction** - Includes subcategories: learning by example, experimentation, observation, and by discovery.
      - One of the most active areas, can apply statistics/math
    - **Learning by analogy** - Combines both deductive and inductive learning. (Being bitten by a teased dog may make an individual not tease bees.)

12-28

Modeling Human Intelligence

- Common Sense
  - Problems that seemed to be most difficult, such as playing chess, turned out to be relatively simple.
  - The computer must be able to make inferences from the knowledge base.
    - Answers to problems might not be listed.
    - The computer will need to come up with its own answers!
    - This has been a very difficult area in Artificial Intelligence.
  - Cyc (enCYClopedia) Computer program that exhibits and can apply common sense.
    - Built by hand! Data painstakingly entered by people
    - e-Cyc: (Electronic commerce) Advanced search engine narrows a search and gives list of meaningful subtopics.

12-29

Neural Networks

- **Neuron**: Basic building-block of the brain.
  - There are several specialized types, but all have the same basic structure:
  - The basic structure of an animal neuron.

12-30

Neural Networks

- Artificial models of the brain are of two distinct types:
  - Electronic: Has electronic circuits that act like neurons.
  - Software: This version runs a program on the computer that simulates the action of the neurons.

12-31

Neural Networks

- **Artificial neurons**: Commonly called processing elements, are modeled after real neurons of humans and other animals.
  - Has many inputs and one output.
    - The inputs are signals that are strengthened or weakened (weighted).
    - If the sum of all the signals is strong enough, the neuron will put out a signal to the output.

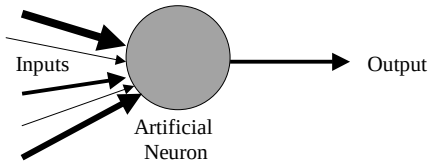

12-32

Neural Networks

- **Neural Network**: A collection of neurons which are interconnected. The output of one connects to several others with different strength connections.
  - Initially, neural networks have no knowledge. (All information is learned from experience using the network.)

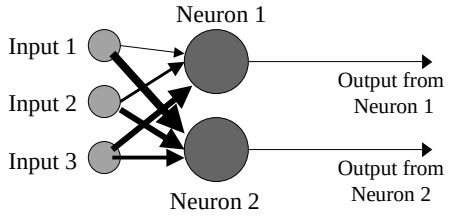

12-33

Neural Networks

- Training a Neural Network
  - **Supervised training**:
    - Occurs when the neural network is given input data.
    - The resulting output is compared to the correct input.
    - The strengths of the connections are then modified so as to minimize errors in succeeding input/output pairs.
  - Example: Back propagation: This method of learning is divided into two phases:
    1. The inputs are applied to the network, and the outputs compared with the correct output.
    2. The resulting information about any error is fed backwards through the network, adjusting the connection strengths to minimize the error.

12-34

Neural Networks

- Neural networks in action: A case study.
  - Mortgage Risk Evaluator.
    - Data from several thousand mortgage applicants was used to train a neural network.
      - Credit data of each individual was paired with each loan result.
      - Patterns for successful loans and defaults of mortgages were contained in the data.
      - The neural network's weights (measurements of strengths) were adjusted to match the actual output.
    - Now, a new mortgage applicant is entered as input. The program determines whether they are a bad risk.
  - Lots of other examples
    - Driving a car, classifying disease, balancing a stick, parsing language

12-35

Evolutionary Systems

- Alan Turing, in 1950, identified three attributes that are the basis for what is now termed genetic programming.
  - Heredity
  - Mutation
  - Natural selection

  - Evolution is being used to create or grow programs.

12-36

Evolutionary Systems

- **Genetic Algorithm** (simulated evolution):
  - Mimics the processes in the genetics of living systems.
  - Created by John Holland (mid-1960's) U. of Michigan.
  - The human puts together the system and specifies the desired results, but the details on how it is done are left to evolve.
  - Example: Koza, a student of Holland, developed a system that had tree-structured chromosomes.
    - Using basic astronomical data, his system came up with Kepler's 3rd law of planetary motion.
      - "the cube of a planet's distance from the sun is proportional to the square of its period"
  - Major problem with genetic algorithms: An intimate knowledge of the system must be known.

12-37

Evolutionary Systems

- **Genetic Programming:**
  - A technique that follows Darwinian evolution.
  - The evolution takes place directly on the programs in the population that are striving to reach the goal specified by the programmer.
    - Only the goal is known and possibly some of the structure of the solution..

12-38

Complex Adaptive Systems

- **Complex adaptive systems**: A collection of many parts individually operating under relatively simple rules, and are highly interactive in a nonlinear way.
  - Their parts are self organizing, operate in parallel, and exhibit emergent behavior (totally unpredictable results can occur).
  - The system of parts evolves with natural selection operating.
  - Example: Mound-building termite colonies in Australia.
    - Mounds can be several feet high.
    - Termites follow a simple set of rules.
    - Mounds affect what can grow around it.

12-39

Complex Adaptive Systems

- **Chaos**:
  - Described as a situation where things seem unorganized and unpredictable.
  - Tiny changes in the starting point produce solutions to a problem that seem to have almost random results.
  - "Butterfly affect": A tiny flip of a butterfly's wings could start a hurricane.
- **Artificial life**: (**a-life**)
  - A phenomena in computers that has attributes of life.
  - Some argue that computer viruses are a form of a-life.
  - A great venue for simulating evolutionary and biological experiments

12-40

Some Requistes for Life

- Autonomy
- Metabolism
- Survival Instinct
- Self-Reproduction
- Evolution
- Adaptation

One can argue that all of these things can be implemented on a computer system

12-41