ShanghaiTech University 上海科技大学 School of Information Science and Technology 信息科学与技术学院

Introduction to Information Science and Technology (EE 100)
Part II: Intelligent Machines and Robotics
Artificial Intelligence - AI

Sören Schwertfeger / 师泽仁

ShanghaiTech University

General Control Scheme for Mobile Robot Systems

Lecture 6:
Artificial Intelligence (AI)

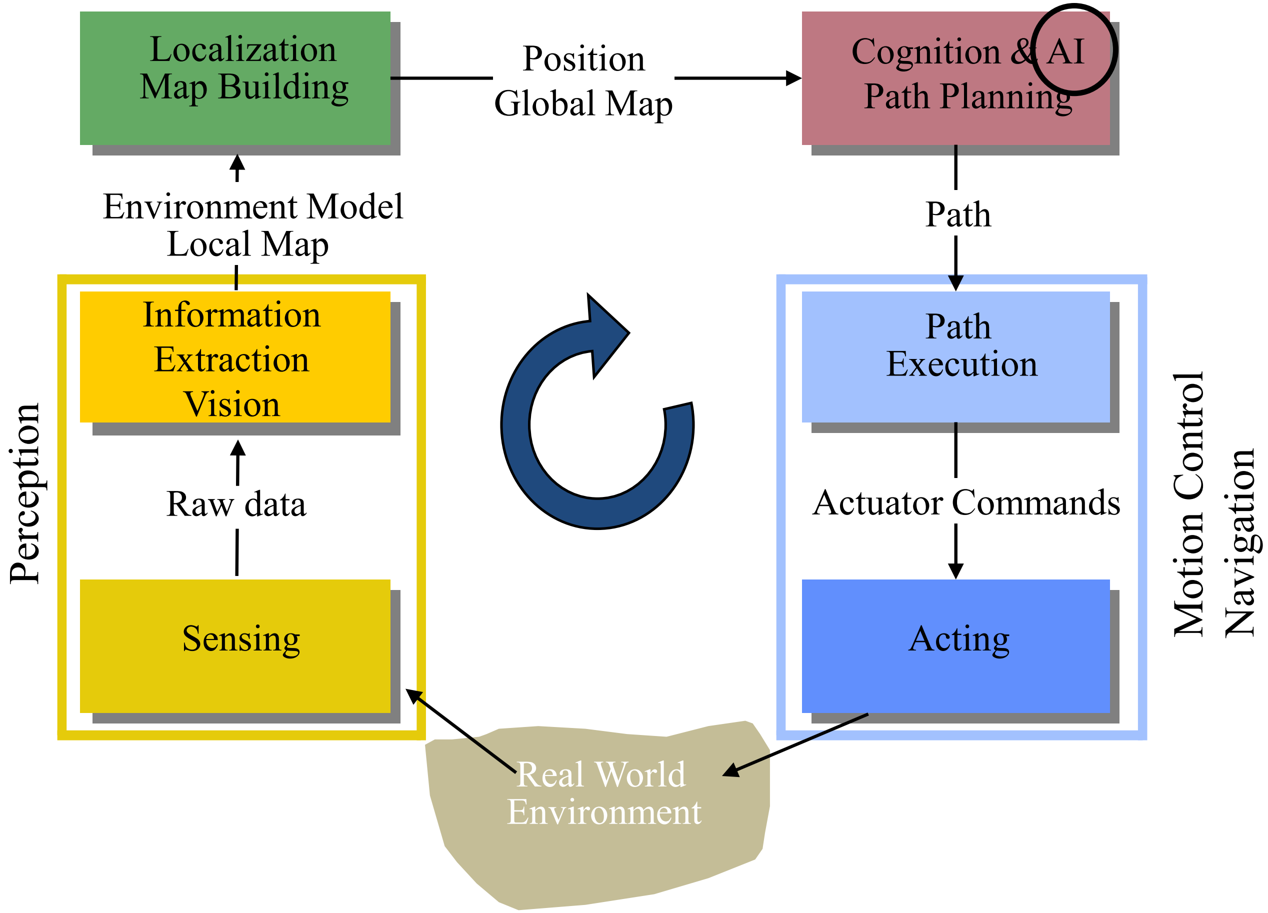

With material from Roland Siegwart and Davide Scaramuzza, ETH Zurich

A BRIEF OVERVIEW OF AI

Artificial Intelligence

- **Artificial intelligence (AI)** is the branch of computer science that develops machines and software with intelligence.
  - AI: Intelligence = rationality
  - Rationality: *maximize goal achievement given available information*

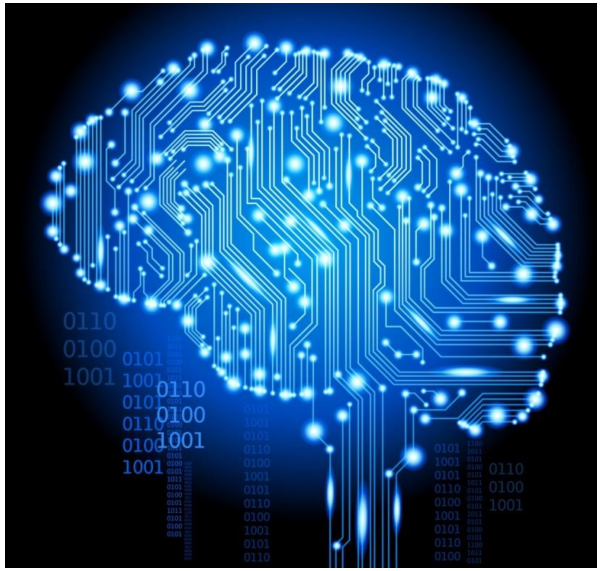

Foundations of AI

- Philosophy
  - Logic, reasoning, mind as a physical system, foundations of learning, language and rationality.
- Mathematics
  - Formal representation and proof algorithms, computation, (un)decidability, (in)tractability, probability.
- Psychology
  - adaptation, phenomena of perception and motor control.
- Economics
  - formal theory of rational decisions, game theory.
- Linguistics
  - knowledge representation, grammar.
- Neuroscience
  - physical substrate for mental activities.
- Control theory
  - homeostatic systems, stability, optimal agent design.

History of AI

- 1943: early beginnings
  - McCulloch & Pitts: Boolean circuit model of brain
- 1950: Turing
  - Turing's "Computing Machinery and Intelligence“
- 1956: Birth of AI
  - Dartmouth meeting: "Artificial Intelligence“ name adopted
- 1950s: initial promise
  - Early AI programs, including
  - Samuel's checkers program
  - Newell & Simon's Logic Theorist
- 1955-65: “great enthusiasm”
  - Newell and Simon: GPS, general problem solver
  - Gelertner: Geometry Theorem Prover
  - McCarthy: invention of LISP

History of AI

- 1966—73: Reality dawns
  - Realization that many AI problems are intractable
  - Limitations of existing neural network methods identified
    - Neural network research almost disappears
- 1969—85: Adding domain knowledge
  - Development of knowledge-based systems
  - Success of rule-based expert systems,
    - Very narrow field; did not scale well in practice
- 1986-- Rise of machine learning
  - Neural networks return to popularity
  - Major advances in machine learning algorithms and applications
- 1990-- Role of uncertainty
  - Bayesian networks as a knowledge representation framework
- 1995-- AI as Science
  - Integration of learning, reasoning, knowledge representation
  - AI methods used in vision, language, data mining, etc

HAL: from the movie 2001

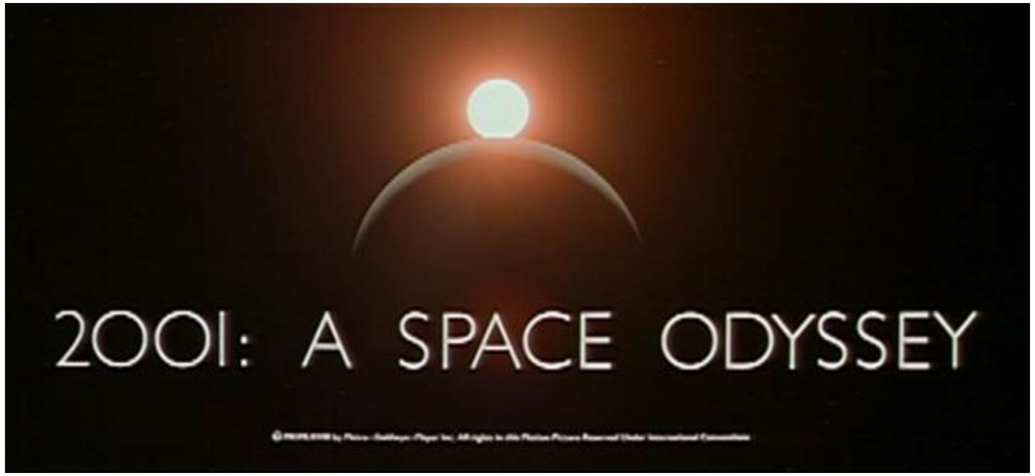

- *2001: A Space Odyssey*
  - classic science fiction movie from 1969
- HAL
  - part of the story centers around an intelligent computer called HAL
  - HAL is the “brains” of an intelligent spaceship
  - in the movie, HAL can
    - speak easily with the crew
    - see and understand the emotions of the crew
    - navigate the ship automatically
    - diagnose on-board problems
    - make life-and-death decisions
    - display emotions

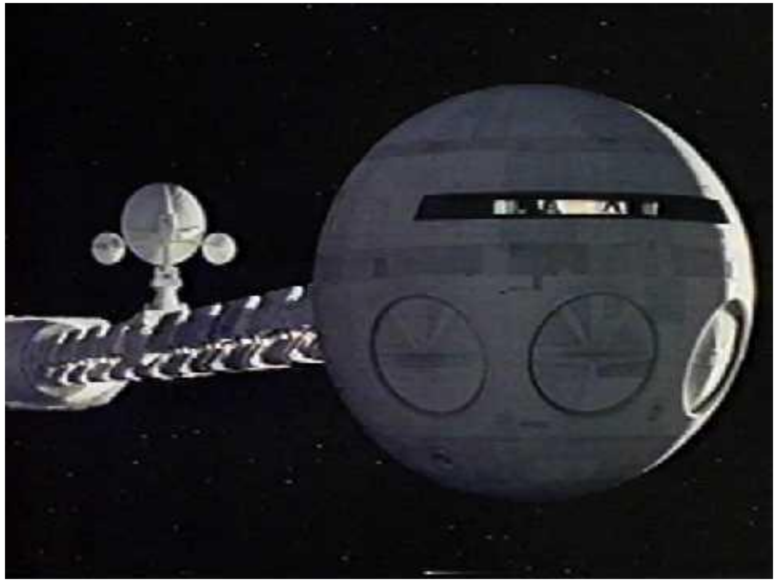

- In 1969 this was science fiction: is it still science fiction?

Consider what might be involved in building a computer like Hal….

- What are the components that might be useful?
  - Fast hardware?
  - Chess-playing at grandmaster level?
  - Speech interaction?
    - speech synthesis
    - speech recognition
    - speech understanding
  - Image recognition and understanding ?
  - Learning?
  - Planning and decision-making?

Can we build hardware as complex as the brain?

- How complicated is our brain?
  - a neuron, or nerve cell, is the basic information processing unit
  - estimated to be on the order of 10^{12} neurons in a human brain
  - many more synapses (10^{14}) connecting these neurons
  - cycle time: 10^{-3} seconds (1 millisecond)
- How complex can we make computers?
  - 10^8 or more transistors per CPU
  - supercomputer: hundreds of CPUs, 10^{12} bits of RAM
  - cycle times: order of 10^{-9} seconds
- Conclusion
  - YES: in the near future we can have computers with as many basic processing elements as our brain, but with
    - far fewer interconnections (wires or synapses) than the brain
    - much faster updates than the brain
  - but building hardware is very different from making a computer behave like a brain!

Can Computers beat Humans at Chess?

- Chess Playing is a classic AI problem
  - well-defined problem
  - very complex: difficult for humans to play well

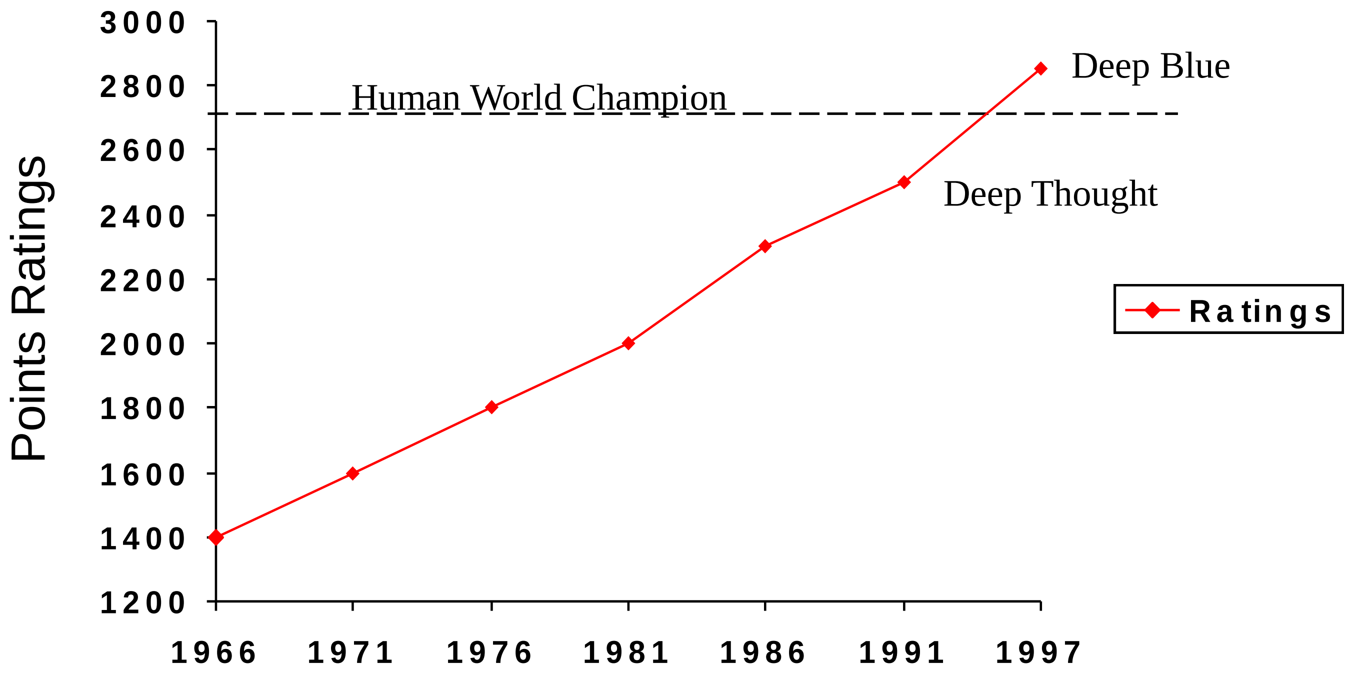

- Conclusion:
  - YES: today’s computers can beat even the best human

Can Computers Recognize Speech?

- Speech Recognition:
  - mapping sounds from a microphone into a list of words
  - classic problem in AI, quite difficult

- Recognizing single words from a small vocabulary
  - systems can do this with high accuracy (order of 99%)
  - e.g., directory inquiries
    - limited vocabulary (area codes, city names)
    - computer tries to recognize you first, if unsuccessful hands you over to a human operator
    - saves millions of dollars a year for the phone companies

Recognizing human speech (ctd.)

- Recognizing normal speech is much more difficult
  - speech is continuous: where are the boundaries between words?
    - e.g., “John’s car has a flat tire”
  - large vocabularies
    - can be many thousands of possible words
    - we can use **context** to help figure out what someone said
      - e.g., hypothesize and test
      - try telling a waiter in a restaurant:
        “I would like some dream and sugar in my coffee”
  - background noise, other speakers, accents, colds, etc
  - Modern systems become more and more accurate

- Conclusion:
  - NO, normal speech is too complex to accurately recognize – but we are getting better (Siri)
  - YES, for restricted problems (small vocabulary, single speaker)

Can Computers Understand speech?

- Understanding is different to recognition:
  - “Time flies like an arrow”
    - assume the computer can recognize all the words
    - how many different interpretations are there?

Can Computers Understand speech?

- Understanding is different to recognition:
  - “Time flies like an arrow”
    - assume the computer can recognize all the words
    - how many different interpretations are there?
      - 1. time passes quickly like an arrow?
      - 2. command: time the flies the way an arrow times the flies
      - 3. command: only time those flies which are like an arrow
      - 4. “time-flies” are fond of arrows

Can Computers Understand speech?

- Understanding is different to recognition:
  - “Time flies like an arrow”
    - assume the computer can recognize all the words
    - how many different interpretations are there?
      - 1. time passes quickly like an arrow?
      - 2. command: time the flies the way an arrow times the flies
      - 3. command: only time those flies which are like an arrow
      - 4. “time-flies” are fond of arrows
    - only 1. makes any sense,
      - but how could a computer figure this out?
      - clearly humans use a lot of implicit commonsense knowledge in communication

- Conclusion: NO, much of what we say is beyond the capabilities of a computer to understand at present

Can Computers Learn and Adapt ?

- Learning and Adaptation
  - consider a computer learning to drive on the freeway
  - we could teach it lots of rules about what to do
  - or we could let it drive and steer it back on course when it heads for the embankment
    - systems like this are under development (e.g., Daimler Benz)

  - **machine learning** allows computers to learn to do things without explicit programming
  - many successful applications:
    - requires some “set-up”: does not mean your PC can learn to forecast the stock market or become a brain surgeon

- Conclusion: YES, computers can learn and adapt, when presented with information in the appropriate way

Can Computers "see"?

- Recognition v. Understanding (like Speech)
  - Recognition and Understanding of Objects in a scene
    - look around this room
    - you can effortlessly recognize objects
    - human brain can map 2d visual image to 3d "map"
- Why is visual recognition a hard problem?

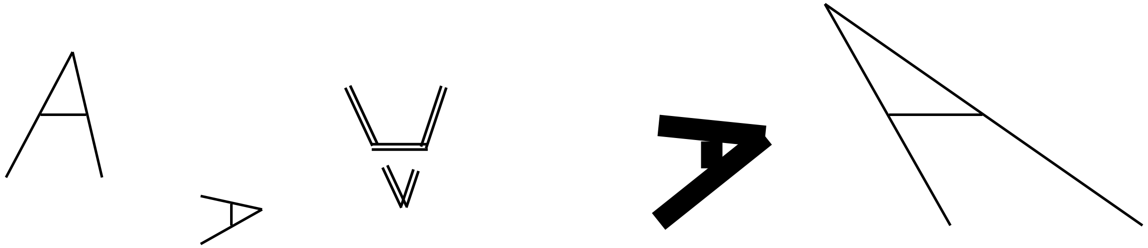

- Conclusion:
  - mostly NO: computers can only "see" certain types of objects under limited circumstances – but we are getting better!
  - YES for certain constrained problems (e.g., face recognition)

Summary of State of AI Systems in Practice

- Speech synthesis, recognition and understanding
  - very useful for limited vocabulary applications
  - unconstrained speech understanding is still too hard
- Computer vision
  - works for constrained problems (hand-written zip-codes)
  - understanding real-world, natural scenes is still too hard
- Learning
  - adaptive systems are used in many applications: have their limits
- Planning and Reasoning
  - only works for constrained problems: e.g., chess
  - real-world is too complex for general systems
- Overall:
  - many components of intelligent systems are “doable”
  - there are many interesting research problems remaining

Acting humanly: Turing test

- Turing (1950) "Computing machinery and intelligence“
- "Can machines think?" → "Can machines behave intelligently?“
- Operational test for intelligent behavior: the Imitation Game

- Suggests major components required for AI:
  - knowledge representation
  - reasoning,
  - language/image understanding,
  - learning

* Question: is it important that an intelligent system act like a human?

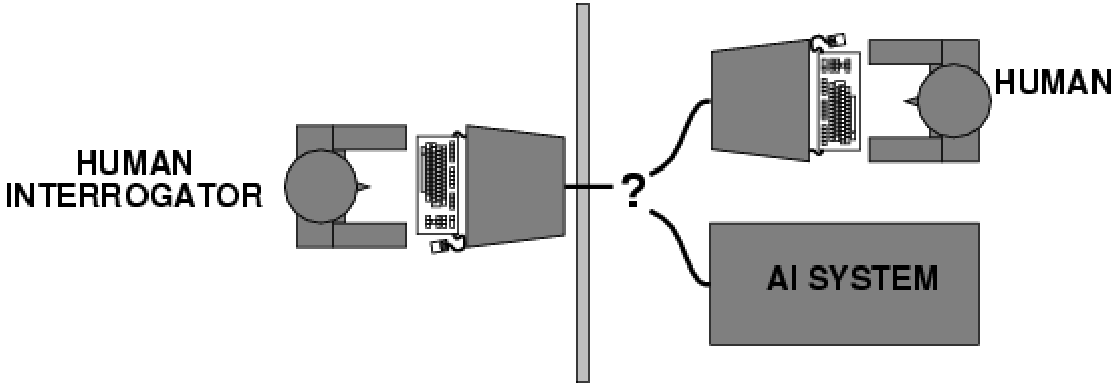

Thinking humanly

- Cognitive Science approach
  - Try to get “inside” our minds
  - E.g., conduct experiments with people to try to “reverse-engineer” how we reason, learning, remember, predict

- Problems
  - Humans don’t behave rationally
    - e.g., insurance, fasten seatbelts
  - The reverse engineering is very hard to do
  - The brain’s hardware is very different to a computer program

Thinking rationally

- Represent facts about the world via logic
- Use logical inference as a basis for reasoning about these facts
- Can be a very useful approach to AI
  - E.g., theorem-provers
- Limitations
  - Does not account for an agent’s uncertainty about the world
    - E.g., difficult to couple to vision or speech systems

Acting rationally

- Decision theory/ Economics
  - Set of future states of the world
  - Set of possible actions an agent can take
  - Utility = gain to an agent for each action/state pair

  - An agent acts rationally if it selects the action that maximizes its “utility”
    - Or expected utility if there is uncertainty

- Emphasis is on autonomous agents that behave rationally (make the best predictions, take the best actions)
  - on average over time
  - within computational limitations (“bounded rationality”)

Subfields of AI: Knowledge Representation and Reasoning

- **Ontology:** formal naming and definition of:
  - types,
  - properties, and
  - interrelationships

  of entities.

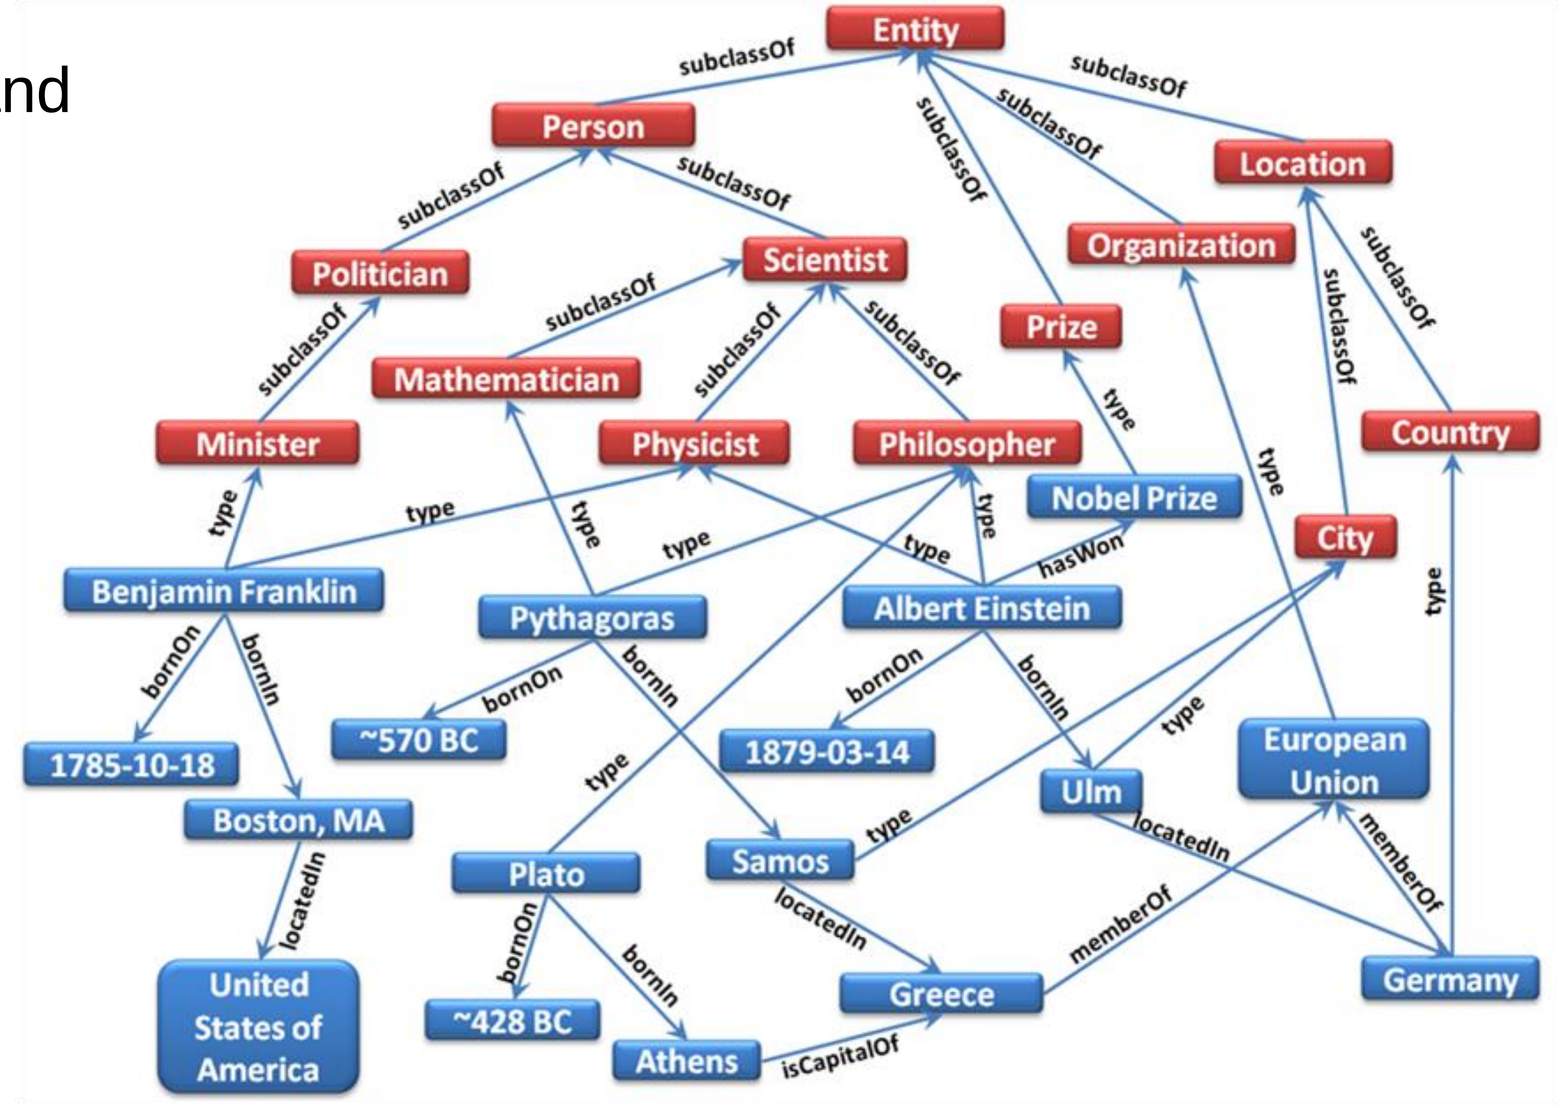

Subfields of AI

- Reasoning: Decision Tree Example
- Give credit to applicant?

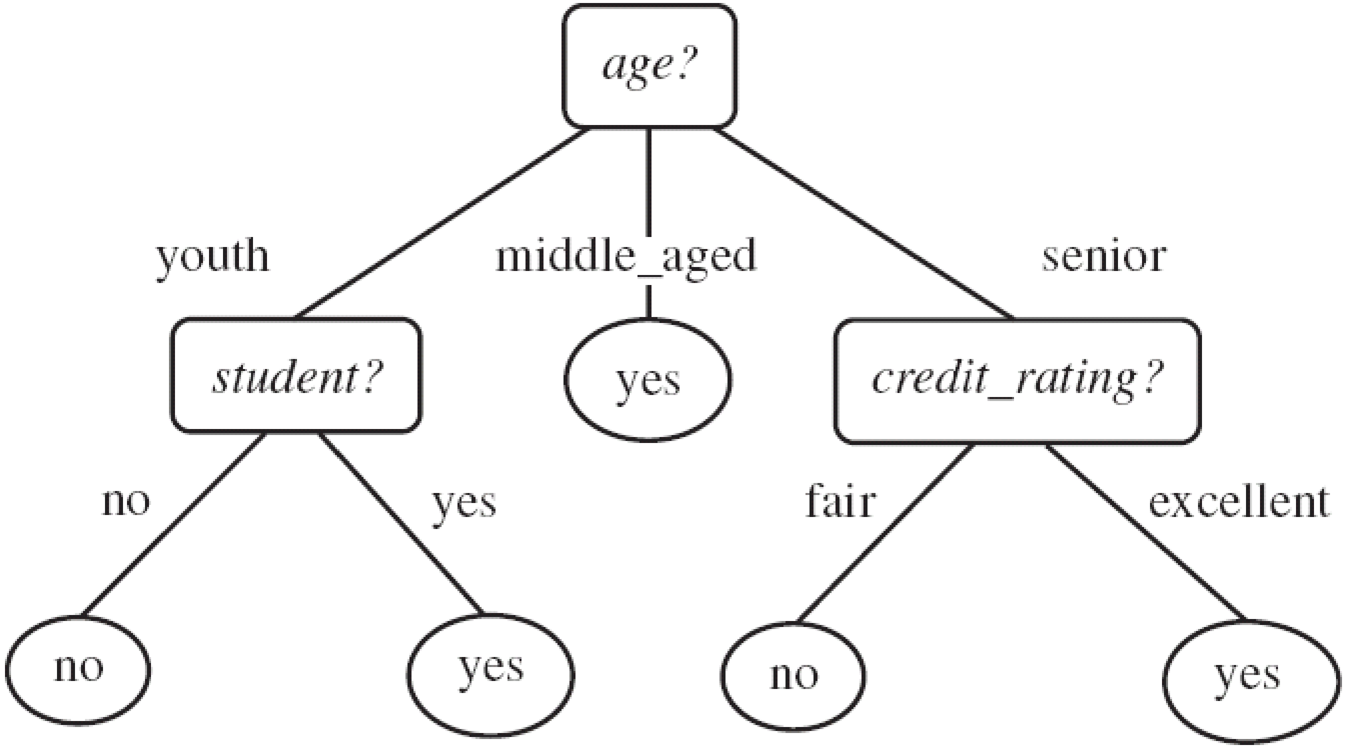

Subfields of AI

- Planning

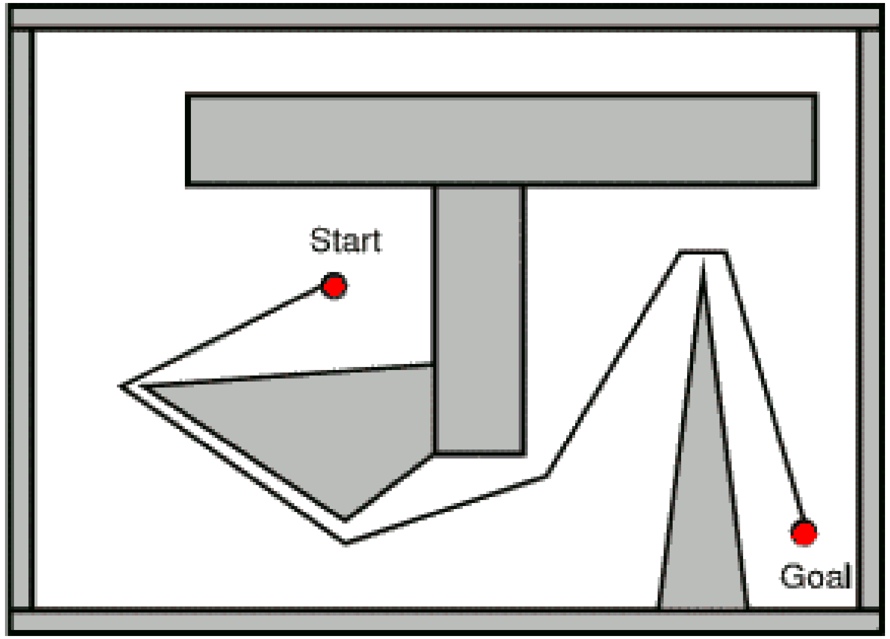

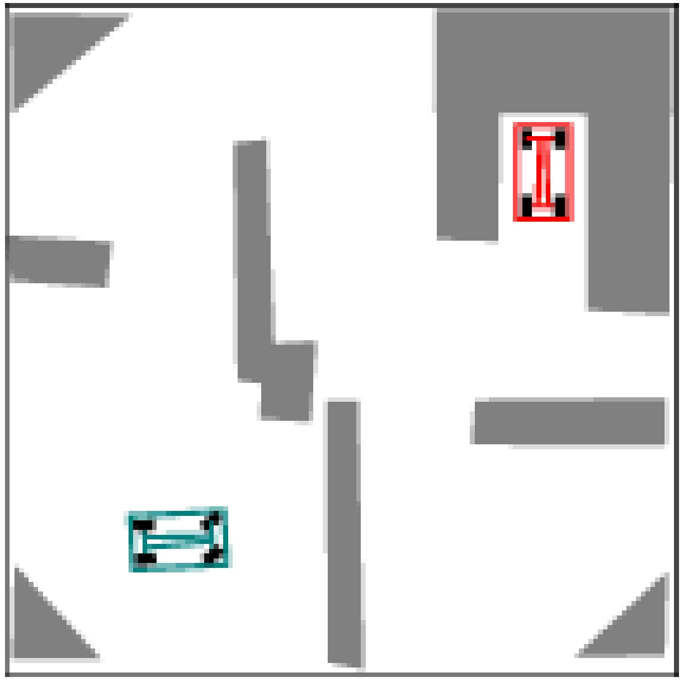

Subfields of AI

- Machine Learning

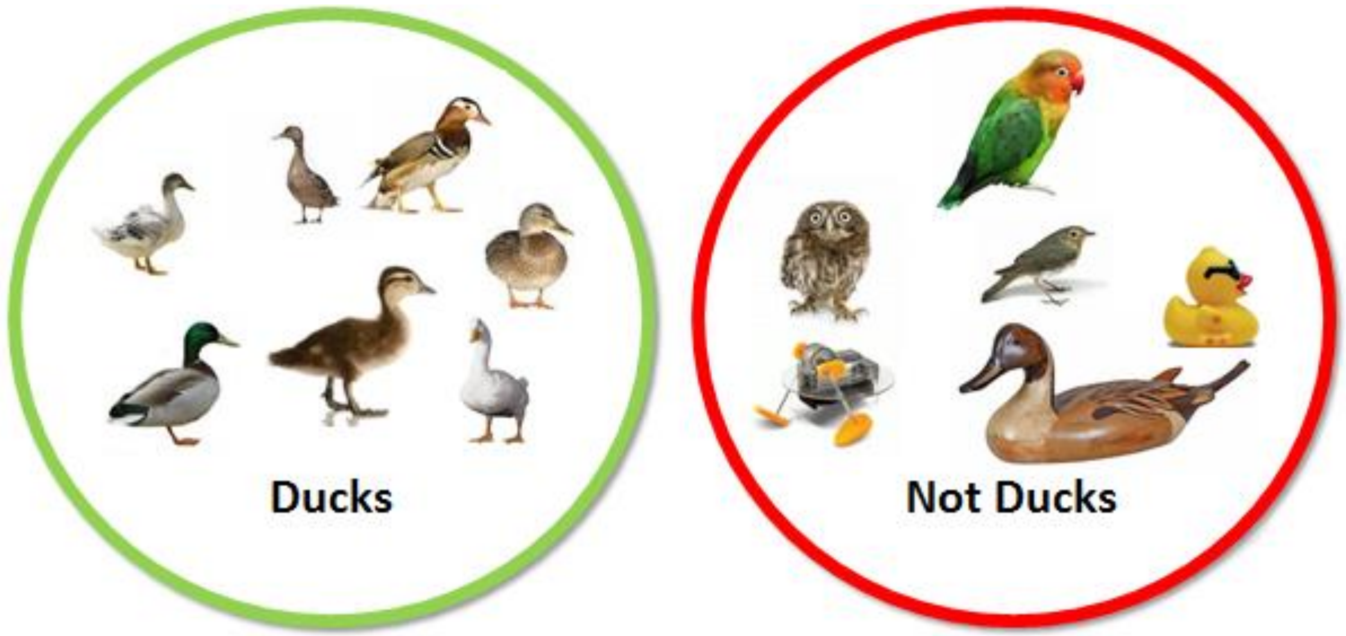

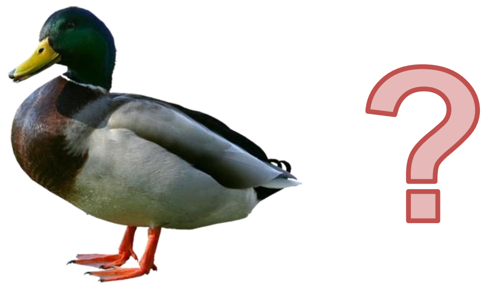

Subfields of AI

- Natural Language Processing

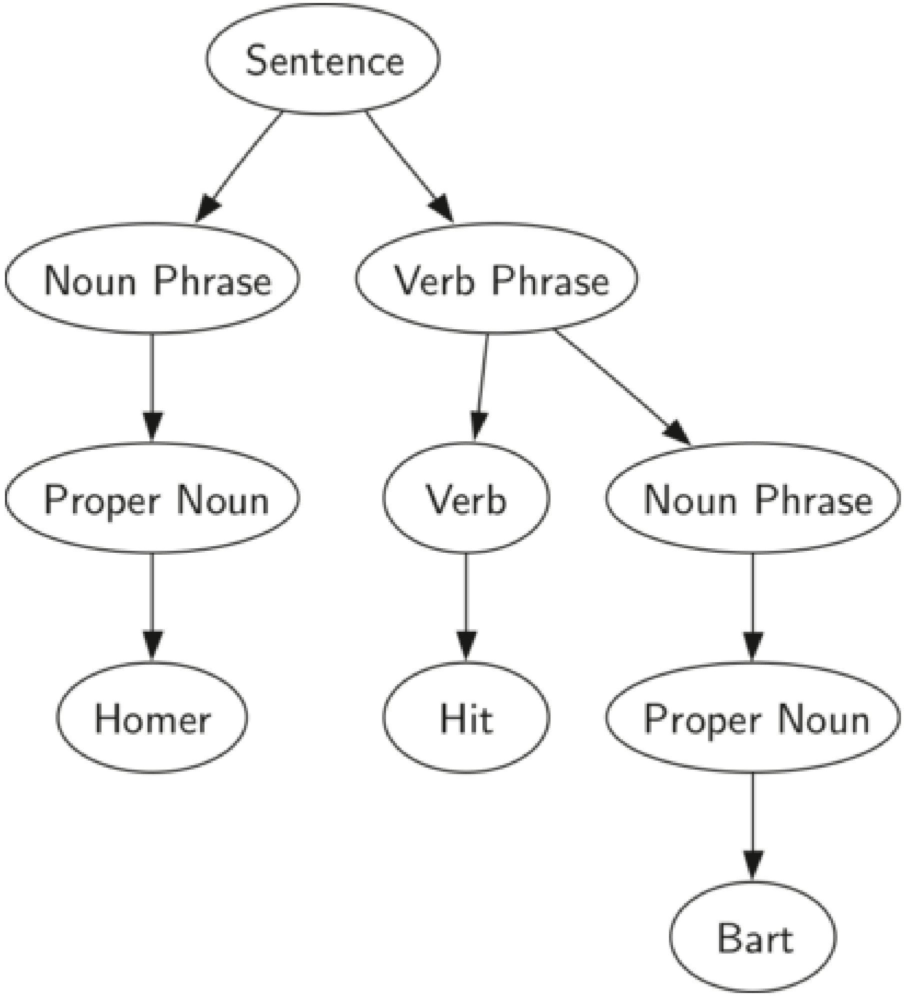

Subfields of AI

- Computer Vision

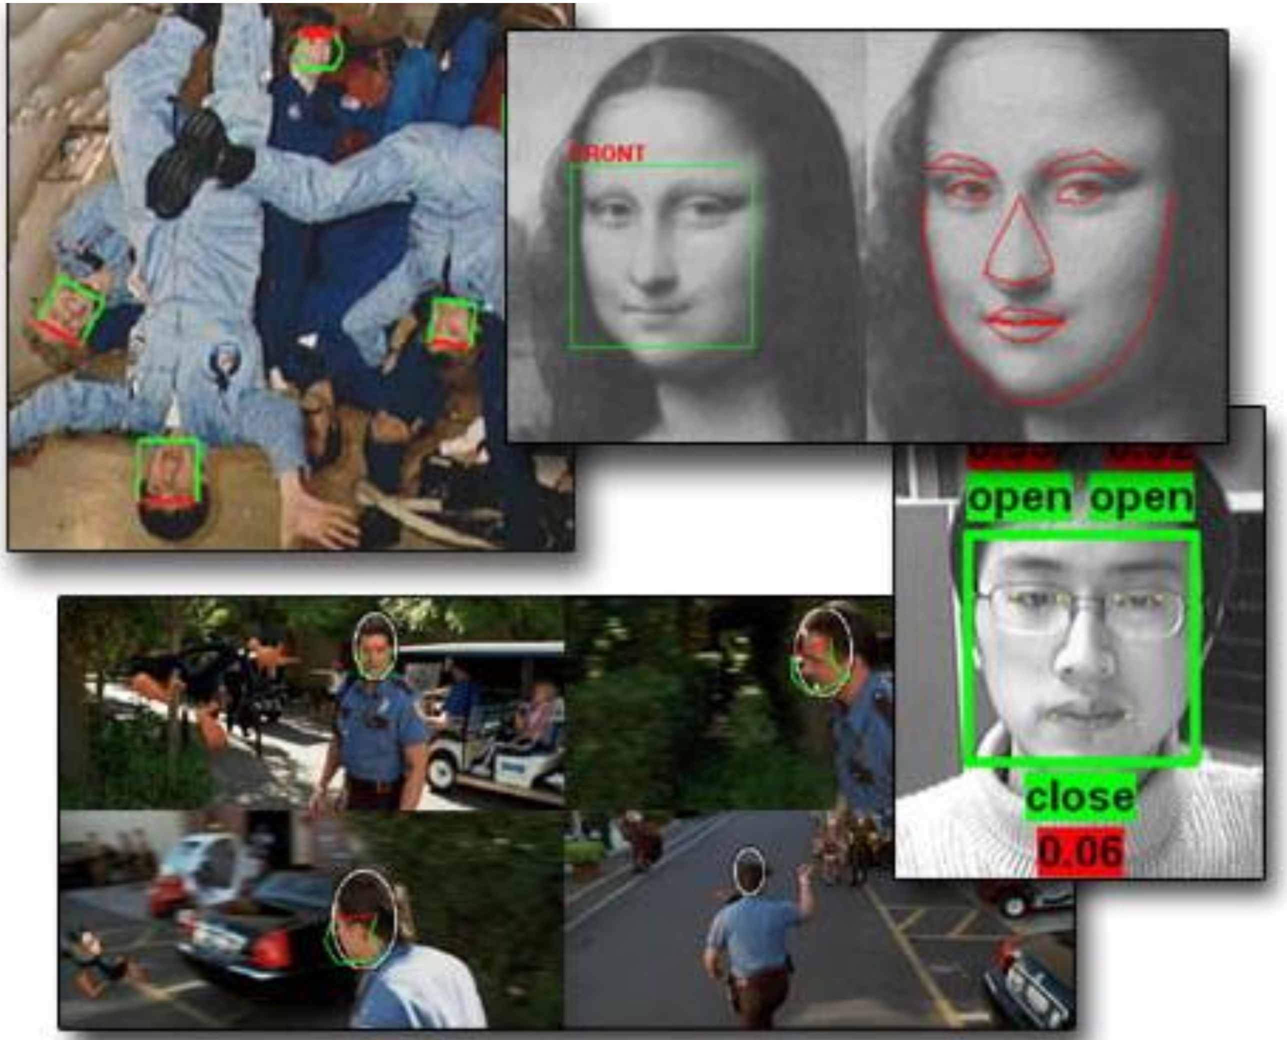

Applications of AI

- Financial trading
- Robots in industrial manufacturing
- Game AI
- Search engine

Applications of AI

- Spam email filter
- Handwriting and speech recognition
- Modern Chinese IME
- Advertisement display on webpages

Applications of AI

- Post Office
  - automatic address recognition and sorting of mail
- Banks
  - automatic check readers, signature verification systems
  - automated loan application classification
- Digital Cameras
  - Automated face detection and focusing
- Customer Service
  - automatic voice recognition

Applications of AI: Machine Translation

- Language problems in international business
  - e.g., at a meeting of Japanese, Korean, Vietnamese and Swedish investors, no common language
  - or: you are shipping your software manuals to 127 countries
  - solution; hire translators to translate
  - would be much cheaper if a machine could do this

- How hard is automated translation
  - very difficult! e.g., English to Russian
    - “The spirit is willing but the flesh is weak” (English)
    - “the vodka is good but the meat is rotten” (Russian)
  - not only must the words be translated, but their meaning also!

- Nonetheless....
  - commercial systems can do a lot of the work very well (e.g., restricted vocabularies in software documentation)
  - algorithms which combine dictionaries, grammar models, etc.
  - Recent progress using “black-box” machine learning techniques

Strong AI vs. Weak AI

- Weak AI (Applied AI)
  - AI that accomplishes specific tasks

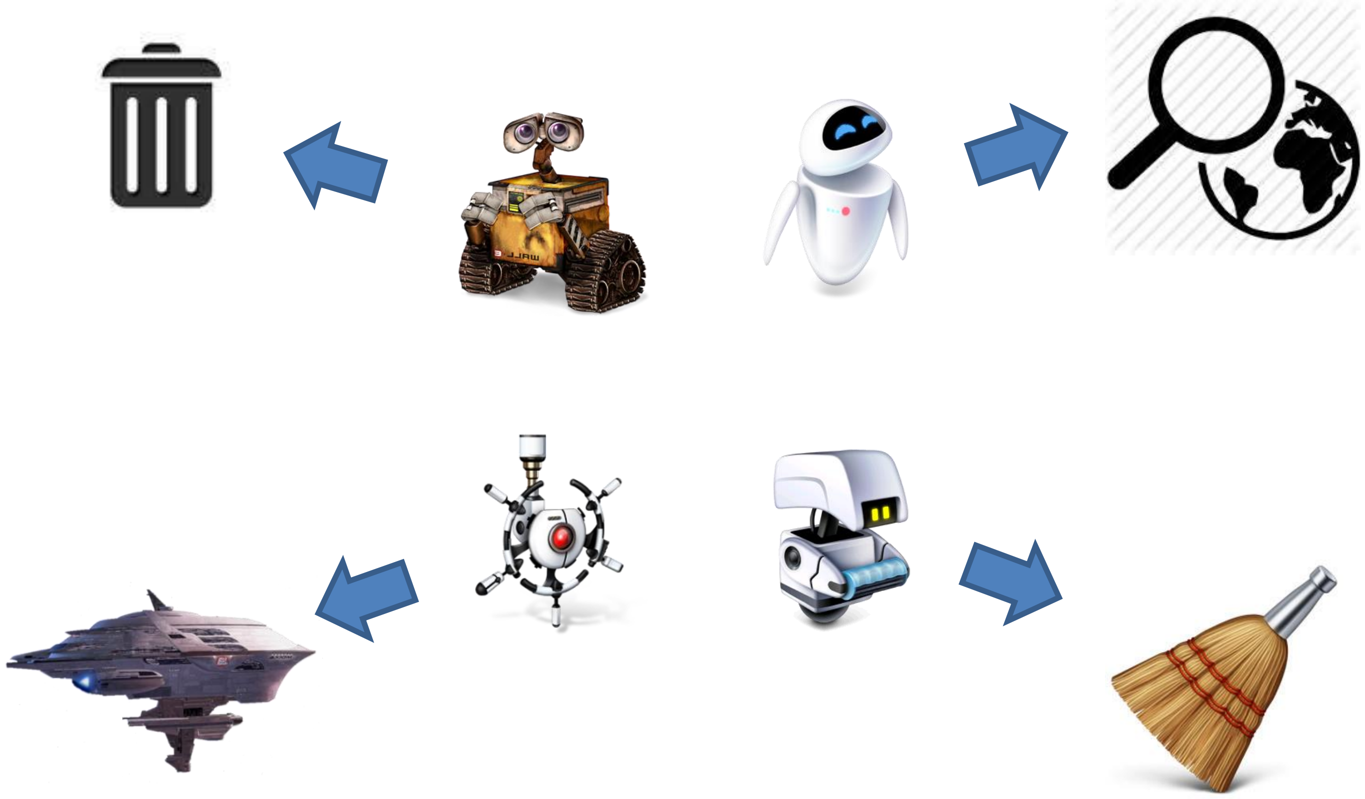

Strong AI vs. Weak AI

- Strong AI (General AI)
  - human-like intelligence – AI that could successfully perform any intellectual task that a human can

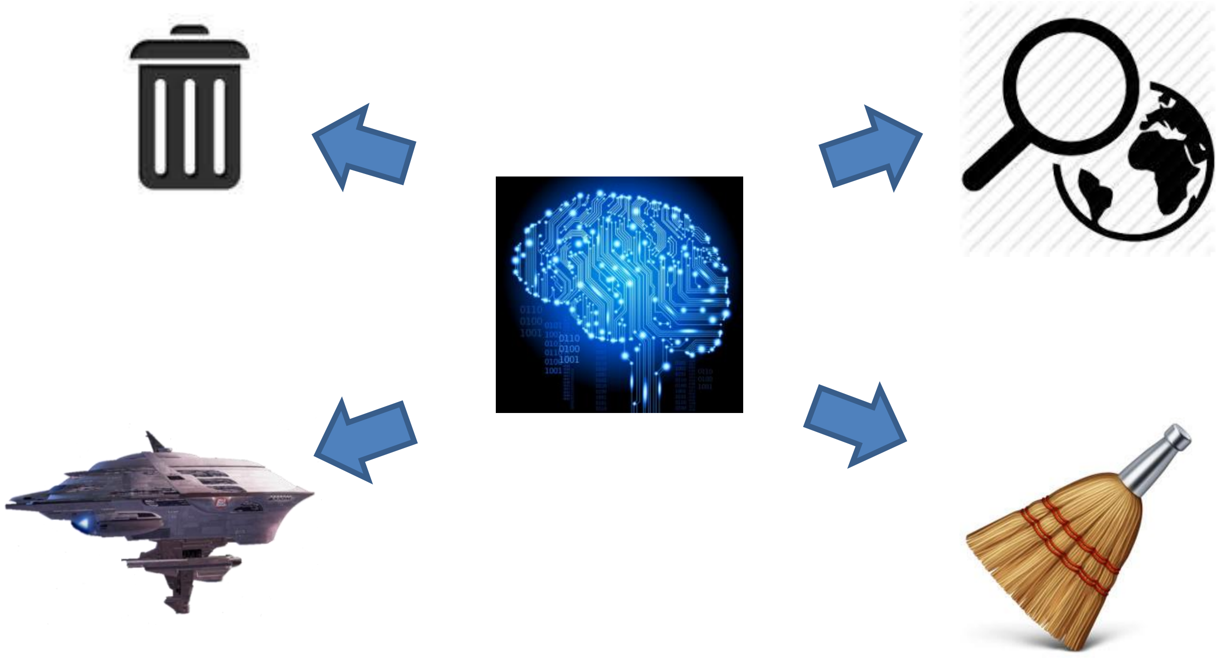

Central problems of strong AI

- Knowledge Representation (KR)
  - Knowledge: facts, beliefs, concepts, skills, … that are accumulated over time

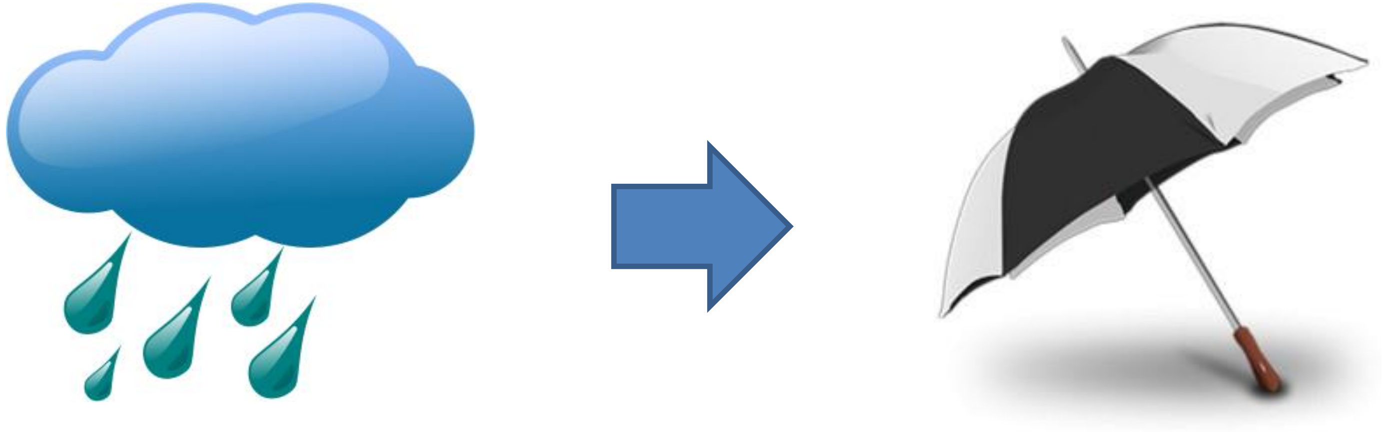

Central problems of strong AI

- Inference
  - How to utilize knowledge to derive new information based on existing information

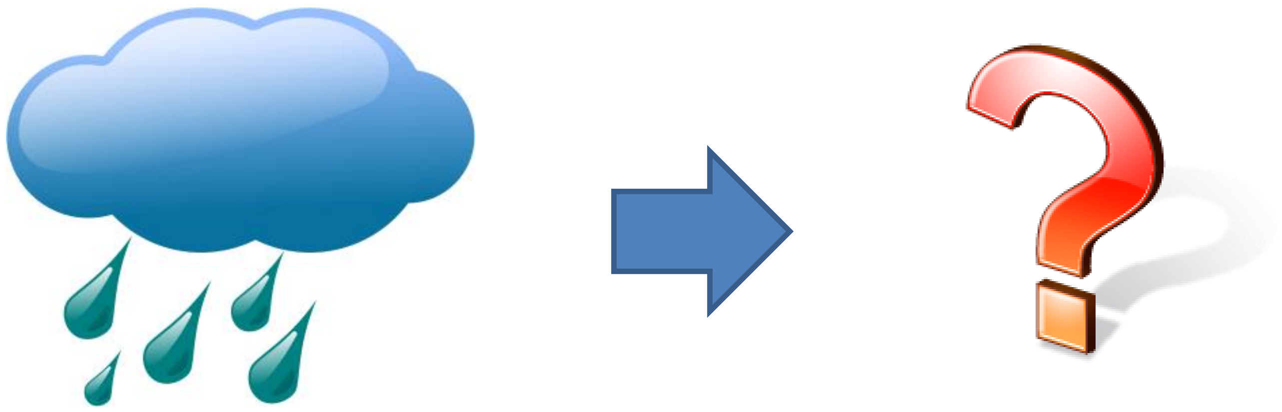

Central problems of strong AI

- Learning
  - How to accumulate knowledge from experience and education

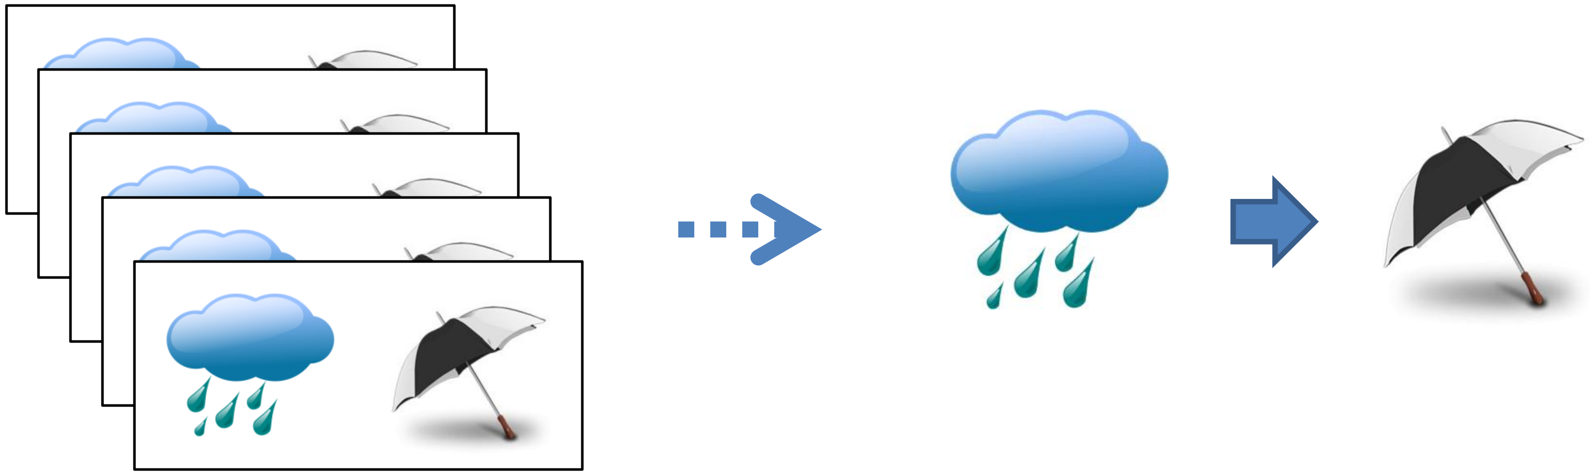

Three types of approaches

- Symbolism
  - Physical symbol system hypothesis

> “A physical symbol system has the necessary and sufficient means for general intelligent action.”
>
> — Allen Newell and Herbert A. Simon (1976)

Three types of approaches

- Symbolism
  - Representing knowledge with symbols and their compositions (expressions)
  - Inference and learning is done by manipulating symbols (e.g., logic)

$$\forall x \forall y, Human(x) \wedge Place(y) \wedge At(x, y) \wedge Rain(y)$$

$$\rightarrow \exists z, Umbrella(z) \wedge Use(x, z)$$

Three types of approaches

- Connectionism
  - Representing knowledge with interconnected networks of simple units
    - Neural networks

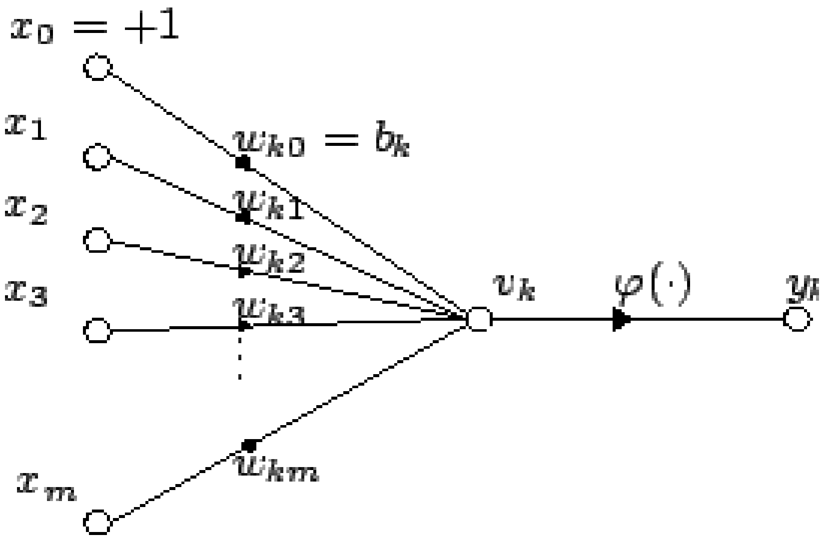

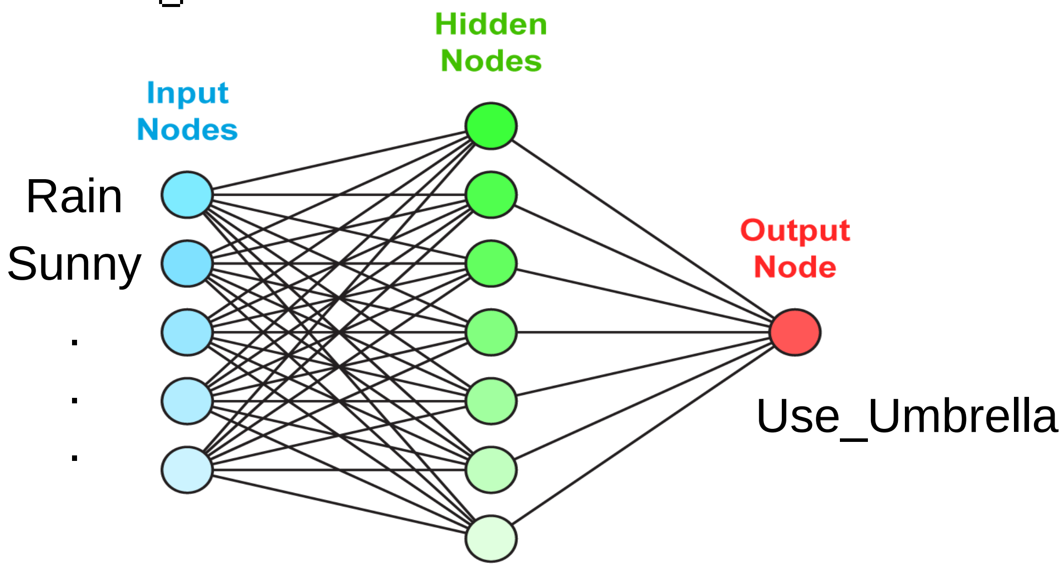

Three types of approaches

- Connectionism
  - Inference
    - computed based on the activation functions of the nodes in the network
  - Learning
    - optimization of connection weights

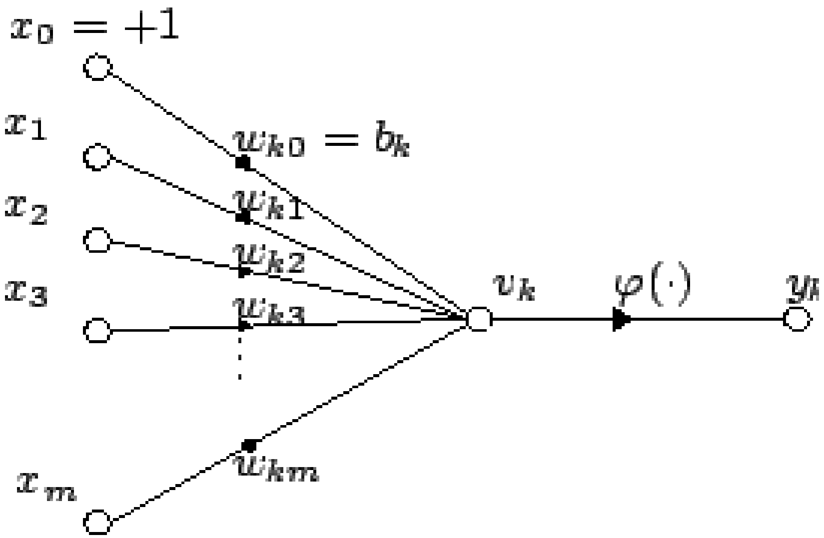

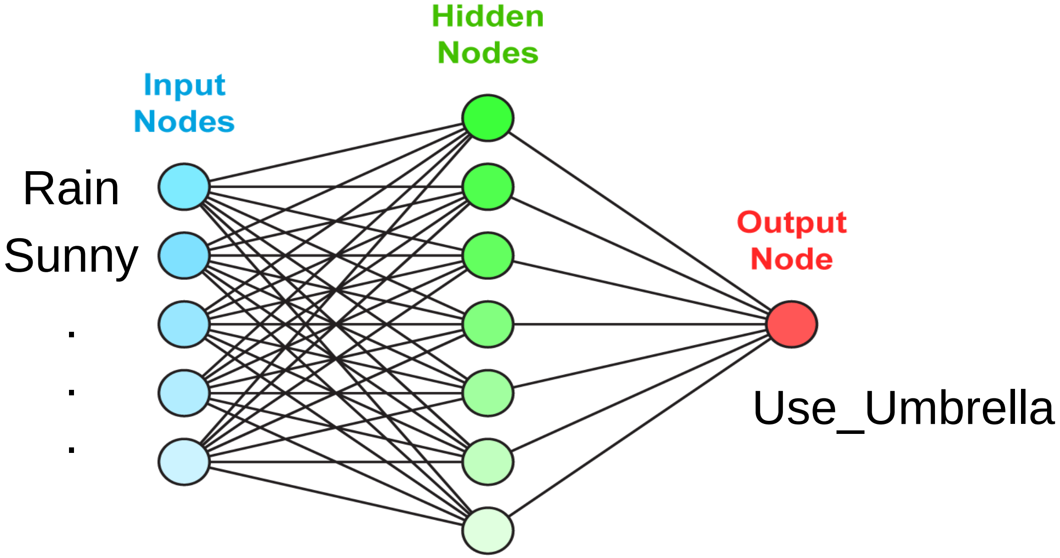

Three types of approaches

- Statistical approaches
  - Representing knowledge with probabilistic models
  - Inference and learning is done by probabilistic inference
  - History: become popular since 1990s

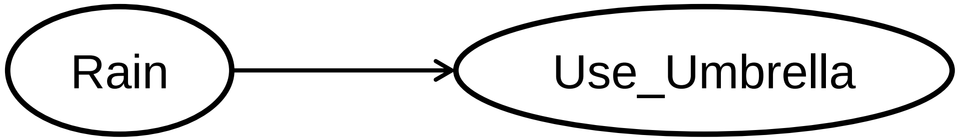

	T	F
	0.2	0.8

Rain	T	F
T	0.95	0.05
F	0.1	0.9

COMPUTERS PLAYING GAMES

Game Playing and Search

- Game playing a long-studied topic in AI
  - Seen as a proxy for how more complex reasoning can be developed

- Search
  - Understand set of possible states
  - Find “best” state or best path to goal state
  - “State” is the condition of the environment
    - e.g. in theorem proving, can be the state of things known
      - By applying known theorems, can expand the state, until reaching the goal theorem

Game Playing

- Abstract AI problem

- Nice and challenging properties
  - Usually state can be clearly, concisely represented
  - Limited number of operations (but can still be large)
  - Unknown factor – account for opponent
  - Search space can be huge
    - Limit response based on time – forces making good “decisions”

Computer Chess

- 2/96: Kasparov vs Deep Blue
  - Kasparov victorious: 3 wins, 2 draws, 1 loss
- 3/97: Kasparov vs Deeper Blue
  - First match won against world champion
  - 512 processors: 200 million chess positions per second

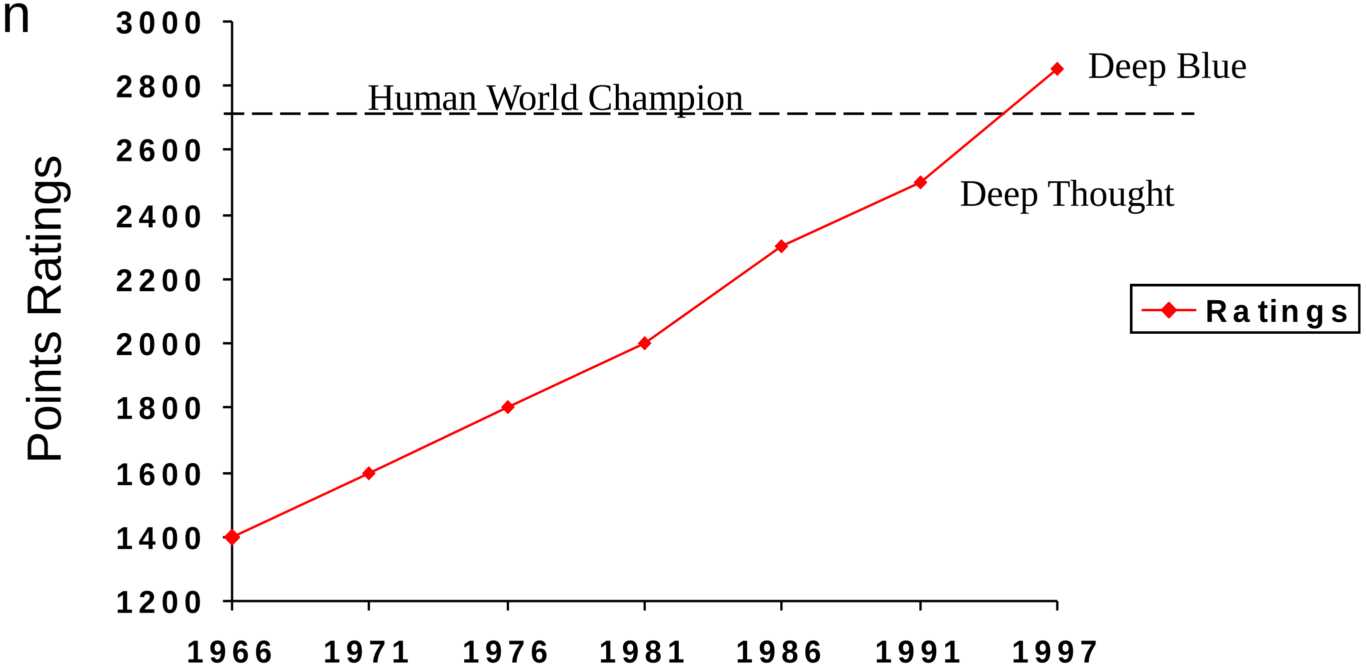

Really Basic State Search Example

- Given a=b, b=c, c=d, prove a=d.

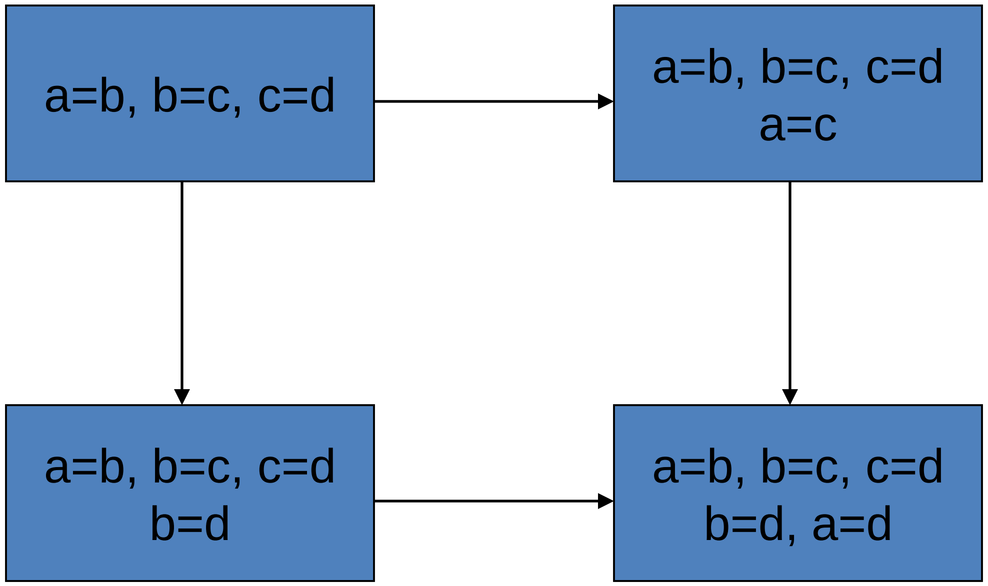

Operators

- Transition from one state to another
  - Fly from one city to another
  - Apply a theorem
  - Move a piece in a game
  - Add person to a meeting schedule

- Operators and states are both usually limited by various rules
  - Can only fly certain routes
  - Only valid moves in game

Search

- Examine possible states, transitions to find goal state
- Interesting problems are those too large to explore exhaustively
- Uninformed search
  - Systematic strategy to explore options
- Informed search
  - Use domain knowledge to limit search => Heuristic

Types of games

- Deterministic vs. random factor
- Known state vs. hidden information

Examples	**Deterministic**	**Chance**
Perfect Info	Chess, Checkers, Othello, Go	Monopoly, Backgammon
Imperfect Info	Stratego, Bridge?	Poker, Scrabble Bridge?

Tick-Tack-Toe

- How to program an AI to play Tick-Tack-Toe?

- 2 players take alternate moves, placing X or O, respectively
- Win if three X or O in one line (horizontal, vertical, diagonal)

- Try all possible moves => exhaustive search
- 3 x 3 = 9 fields => 9 possibilities for first turn
  - 8 fields left for the enemy
  - … 9 * 8 * 7 * 6 * 5 * 4 * 3 * 2 * 1 games = 9! = 362,880
  - Not really:
    - Game ends with 3 in a row
    - Symmetries: Rotated or mirrored states require the same strategy

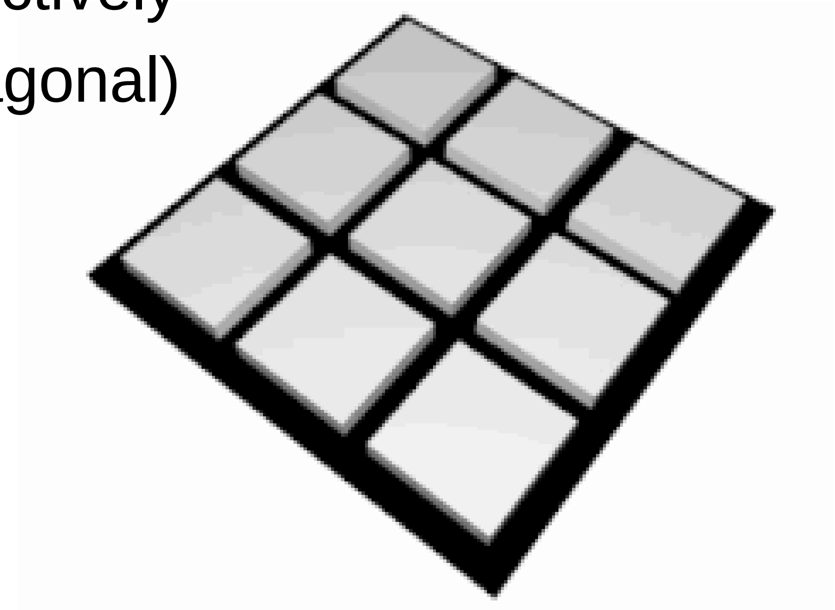

Tick-Tack-Toe

- 3 x 3 = 9 fields => 9 possibilities for first turn
  - Really?
  - Can always rotate or reflect the field => symmetries
  - => 3 possibilities: in the center; on the edge; in the corner:

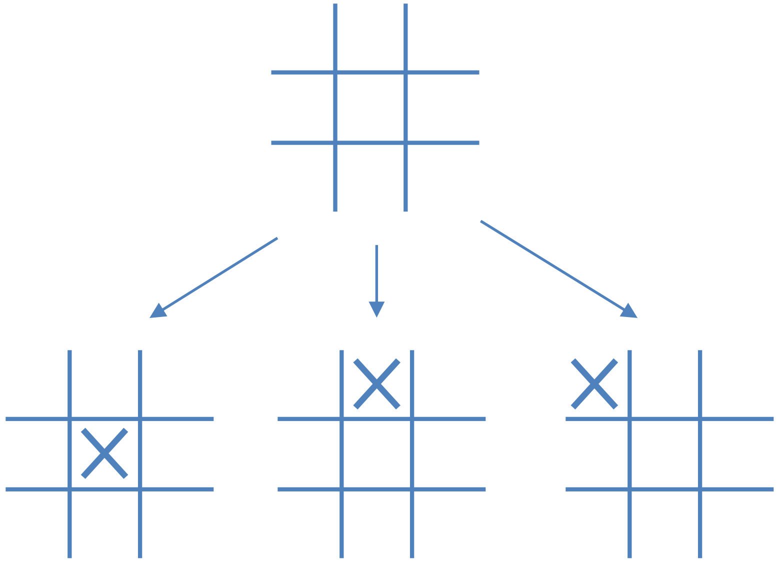

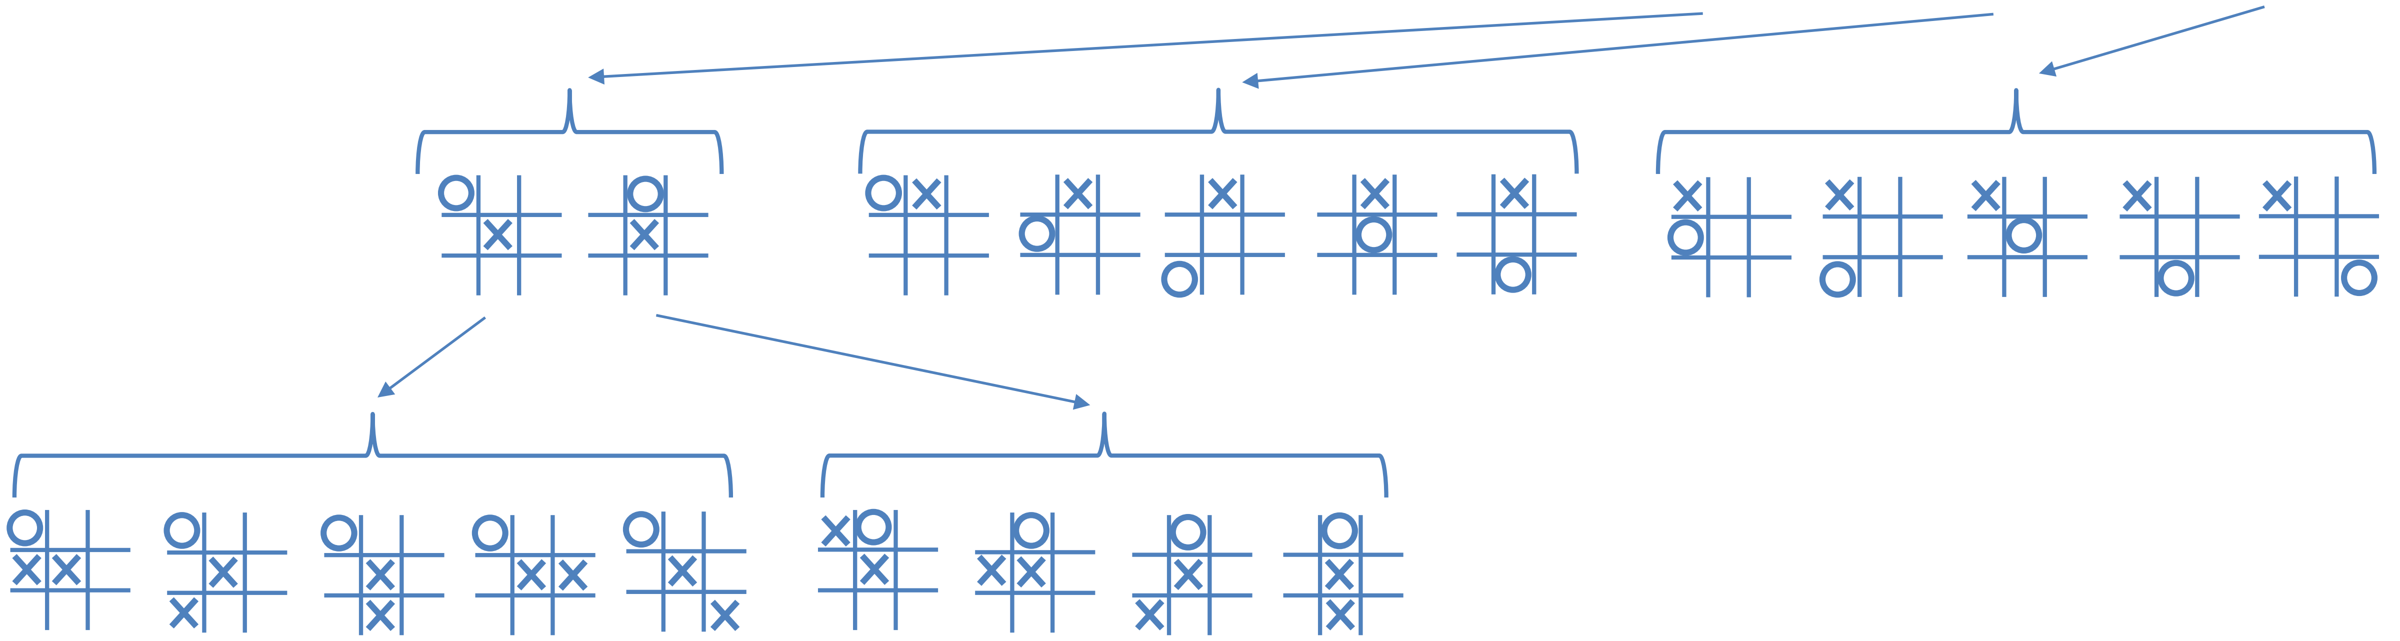

Tick-Tack-Toe

- Only 138 terminal board positions:
    - 91 unique positions are won by (X)
    - 44 unique positions are won by (O)
    - 3 unique positions are drawn

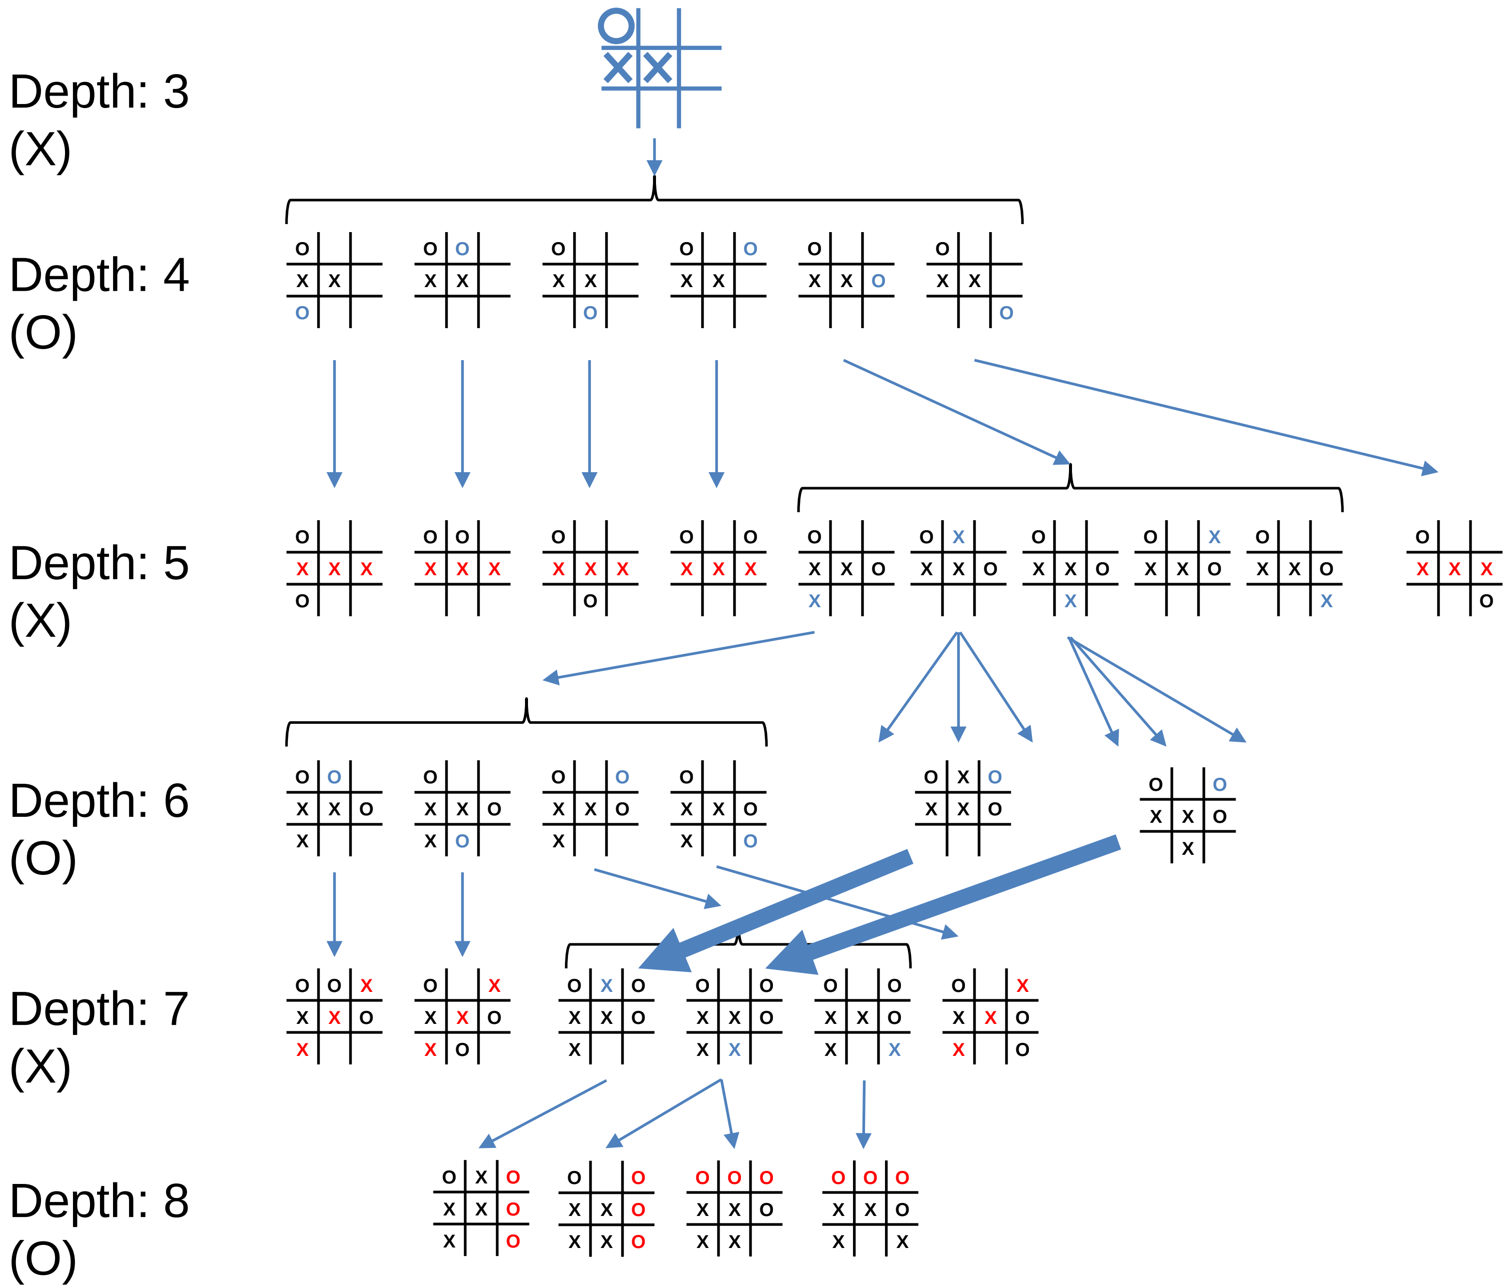

Tick-Tack-Toe

- Two-player zero-sum game:
  - The gain of the one is the loss of the other

- Minimax algorithm:
  - Assume the opponent does what looks best.
  - When it is the human’s turn, pick move that looks best for human;
  - When it is the computer’s turn, pick the move that looks best for the computer
  - *mini*mizing the possible loss for a worst case (*max*imum loss) scenario
  - Associate a value for each position
  - Working backwards from the end of the game

- Tick-Tack-Toe: Whole state space can be computed => optimal strategy (path)

Chess Automaton: Mechanical Turk

- Chess-playing machine: 1770
- By Wolfgang von Kempelen (1734–1804) for Empress Maria Theresa of Austria
- Complex mechanical system
- Strong chess game
- Very high public interest

- Hoax / Lie:
  - Human chess master hidden inside
  - Revealed 1820
  - Amazon Mechanical Turk: Humans perform tasks
    Human Intelligence tasks that AI cannot do – and get paid for it

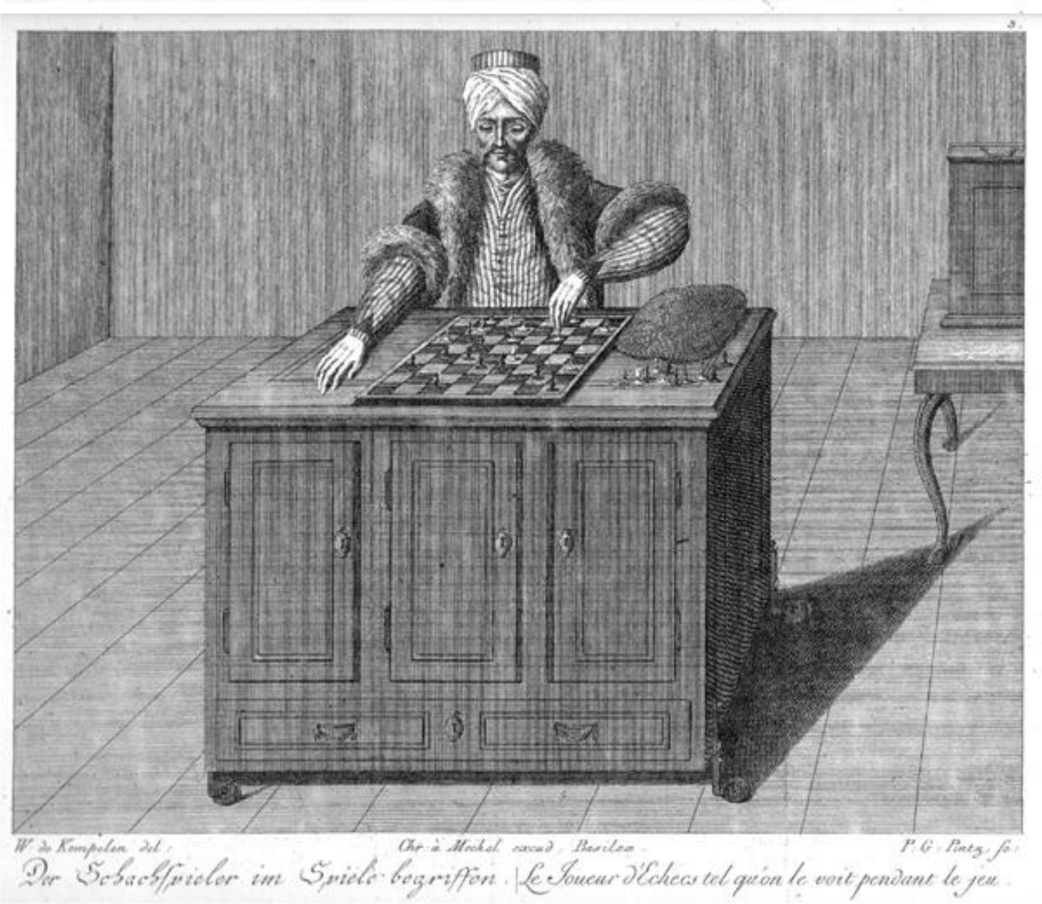

Chess

- Legal positions: 10^{43} - 10^{47}
- Opening: Book of known successful openings
  - Ends with:
    - End of book or
    - Move that is not in book
- Mid Game: Search the tree; use heuristics to evaluate move result:
  - Count and compare sum of value of pieces
  - Development; Control of the center; King safety; Pawn structure; …
- End Game: Very few pieces left => book of all moves possible!
  - Computer plays perfectly – can often win a draw against human even if theoretically lost

Go (围棋)

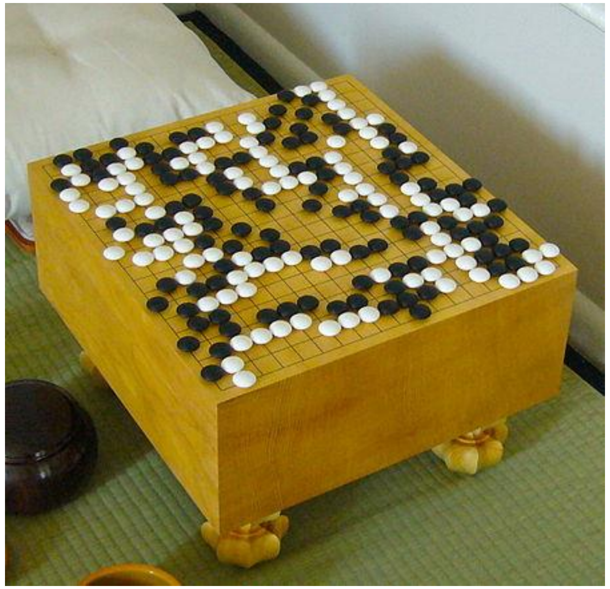

- Number of possible games exceed atoms in the universe (maybe 10^{82})
- Best Go programs only manage to reach amateur dan level
- Go much harder than chess:
  - More field on the board: 361 vs. 64
  - About possible 150–250 moves per turn (always more than 50); chess: about 37
  - Depth 4: 3.2×10^{11} combinations; Depth 8: 5.12×10^{20} (fastest supercomputer would need hours for that)
  - Single stone placement can affect game in 100 moves
  - Heuristic to estimate strength quite hard; Chess much easier (count pieces)

NEXT WEEK:
SIGNALS AND SYSTEMS