

CS5201: Advanced Artificial Intelligence

Introduction



Arijit Mondal

**Dept of Computer Science and Engineering
Indian Institute of Technology Patna**

www.iitp.ac.in/~arijit/

General information

- **Class timings**
 - Monday - 1400-1500
 - Tuesday - 1600-1700
 - Thursday - 1600-1700
- **Venue: LT001**
- **url: <http://www.iitp.ac.in/~arijit/> -> Then visit 'Teaching'**
- **Instructor: Arijit Mondal**
- **TAs**
 - Sophia Jamil
 - Ashish Rao
 - Aditya Singh
 - Jenil Radadiya
 - Anirudh Kandwal
 - Ravi Kumar Paswan
 - Sanjib Chowdhury
 - Shalini Singh
 - Kijen Longren
 - Vikram Kumar
 - Nishant Kumar

Books

- *Artificial Intelligence: A Modern Approach* by Peter Norvig, Stuart Russell
- *Artificial Intelligence - A New Synthesis* by Nils J Nilsson
- *A First Course in Artificial Intelligence* by Deepak Khemani

Evaluation

- Midsem - 30%, Endsem - 50%
- Internal evaluation - 20%
 - At most two quizzes
 - Two projects (in groups, size of group will depend on class strength)
 - Survey on a topic
 - Code implementation
- 75% attendance is mandatory

Introduction

What we are looking for?

- AI is **NOT** just chatGPT or LLM

What we are looking for?

- AI is **NOT** just chatGPT or LLM
- Suppose you need to build a bridge on a river. You asked chatGPT for a design of bridge and necessary details
 - Is it a correct design?
 - How to verify the same?
 - How to provide guarantee
 - Data driven AI / statistical AI is known to have such issues

What we are looking for?

- AI is **NOT** just chatGPT or LLM
- Suppose you need to build a bridge on a river. You asked chatGPT for a design of bridge and necessary details
 - Is it a correct design?
 - How to verify the same?
 - How to provide guarantee
 - Data driven AI / statistical AI is known to have such issues
- Generative AI vs Classical AI

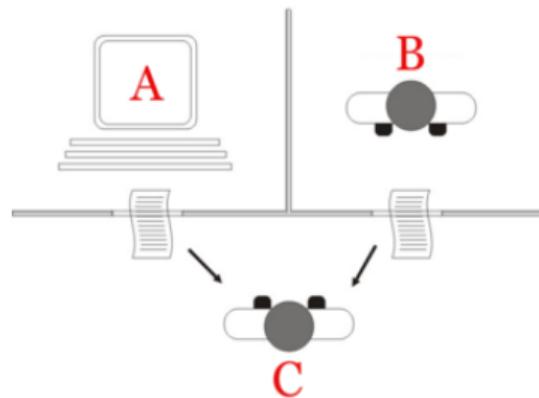
Feature	Generative AI (LLMs)	Classical AI (This Course)
Core Mechanism	Probabilistic Association	Logical Deduction & Search
Best For	Creativity, Writing, Coding, Art	Logistics, Scheduling, Game Theory, Verification
Weakness	Hallucination, Inconsistency	Brittle in messy, unstructured worlds
Explainability	Low (Black Box)	High (Step-by-step Trace)
Goal	Sound Human	Be Correct

What is AI?

- “... the science and engineering of making intelligent machines” (John McCarthy)
- AI is intelligence demonstrated by machines, in contrast to the natural intelligence displayed by humans. (Wikipedia)
- AI attempts not just to understand but also to build intelligent entities. (Russell & Norvig)
- Artificial intelligence is that activity devoted to making machines intelligent, and intelligence is that quality that enables an entity to function appropriately and with foresight in its environment. (Nilsson)

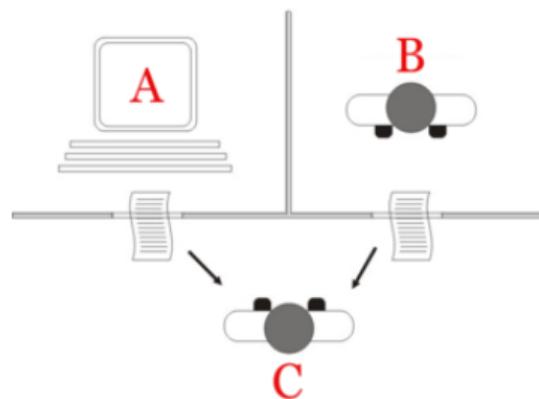
What is AI?

- **Act like human** - The 'Turing test'
 - A human interrogator fails to distinguish responses generated from a computer or a human



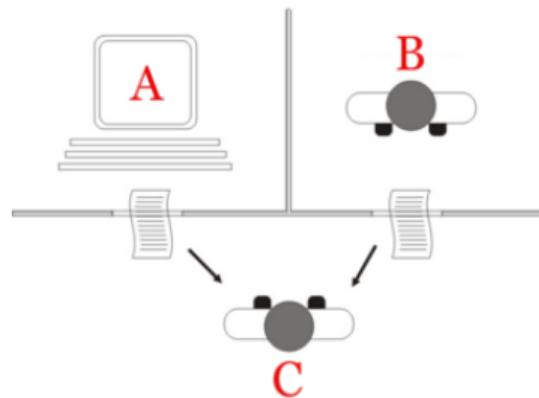
What is AI?

- **Act like human** - The 'Turing test'
 - A human interrogator fails to distinguish responses generated from a computer or a human
- **Think like human** - Cognitive modeling
 - Need to know how human brain functions - through introspection, psychological experiments, brain imaging, etc.



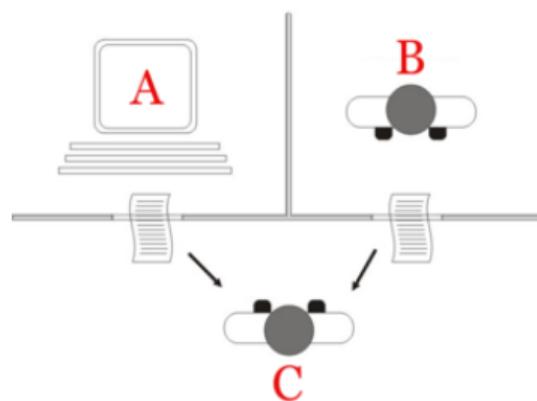
What is AI?

- **Act like human** - The 'Turing test'
 - A human interrogator fails to distinguish responses generated from a computer or a human
- **Think like human** - Cognitive modeling
 - Need to know how human brain functions - through introspection, psychological experiments, brain imaging, etc.
- **Think rationally** - Laws of thought
 - Right thinking, irrefutable reasoning processes



What is AI?

- **Act like human** - The 'Turing test'
 - A human interrogator fails to distinguish responses generated from a computer or a human
- **Think like human** - Cognitive modeling
 - Need to know how human brain functions - through introspection, psychological experiments, brain imaging, etc.
- **Think rationally** - Laws of thought
 - Right thinking, irrefutable reasoning processes
- **Act rationally** - Rational agent
 - Acts to achieve the best outcome in presence of uncertainty, the best expected outcome



What do we prefer?

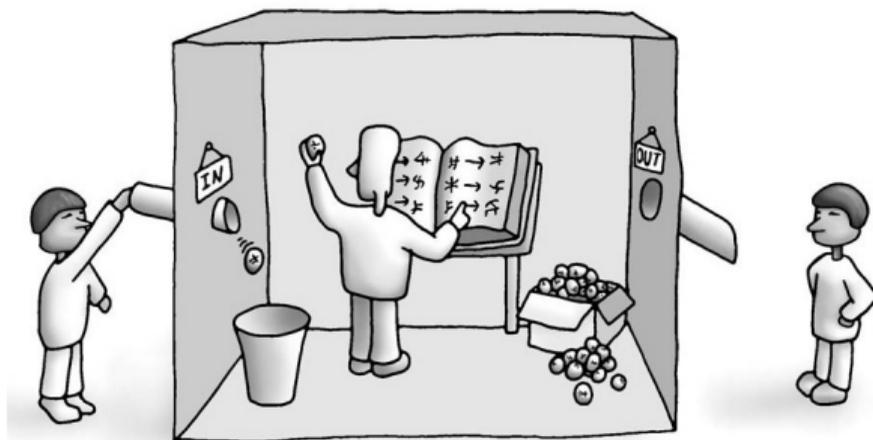
- Humans can make mistakes (we aren't perfectly rational)
 - We want agents that do the best thing, not necessarily the human thing

What do we prefer?

- Humans can make mistakes (we aren't perfectly rational)
 - We want agents that do the best thing, not necessarily the human thing
- If a robot looks both ways, sees nothing, crosses the street, and gets hit by a falling meteorite, was it irrational?
 - It was rational but unlucky

Chinese room argument

- Suppose, artificial intelligence has succeeded in constructing a computer that behaves as if it understands Chinese. It takes Chinese characters as input and, by following the instructions of a computer program, produces other Chinese characters, which it presents as output.
 - Suppose, computer performs its task and comfortably passes the Turing test



Chinese room argument

- Suppose, artificial intelligence has succeeded in constructing a computer that behaves as if it understands Chinese. It takes Chinese characters as input and, by following the instructions of a computer program, produces other Chinese characters, which it presents as output.
 - Suppose, computer performs its task and comfortably passes the Turing test
 - Does the machine literally “understand” Chinese?

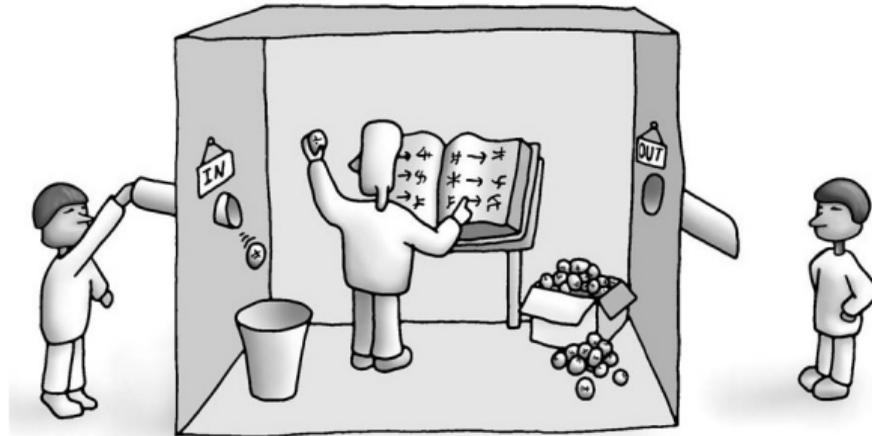
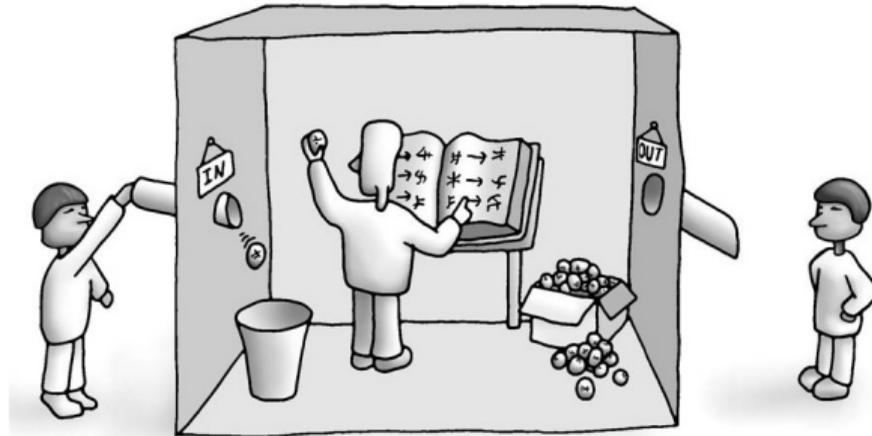


image source: Wikipedia

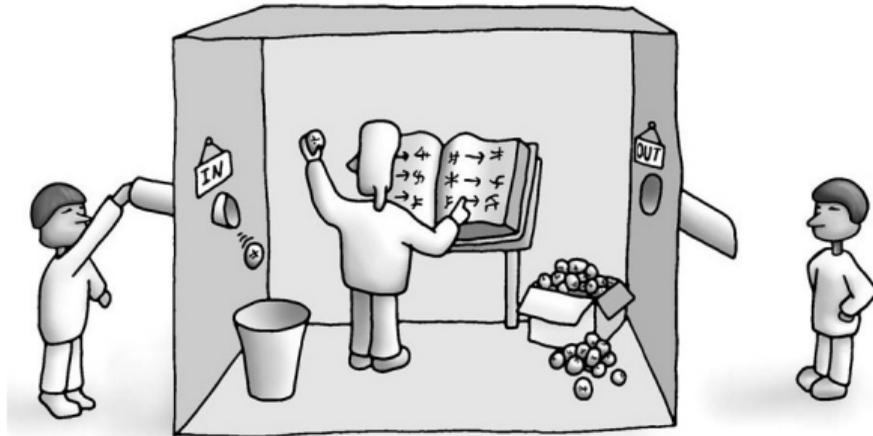
Chinese room argument

- Suppose, artificial intelligence has succeeded in constructing a computer that behaves as if it understands Chinese. It takes Chinese characters as input and, by following the instructions of a computer program, produces other Chinese characters, which it presents as output.
 - Suppose, computer performs its task and comfortably passes the Turing test
 - Does the machine literally “understand” Chinese? – Strong AI



Chinese room argument

- Suppose, artificial intelligence has succeeded in constructing a computer that behaves as if it understands Chinese. It takes Chinese characters as input and, by following the instructions of a computer program, produces other Chinese characters, which it presents as output.
 - Suppose, computer performs its task and comfortably passes the Turing test
 - Does the machine literally “understand” Chinese? – **Strong AI**
 - Is it merely simulating the ability to understand Chinese?



Chinese room argument

- Suppose, artificial intelligence has succeeded in constructing a computer that behaves as if it understands Chinese. It takes Chinese characters as input and, by following the instructions of a computer program, produces other Chinese characters, which it presents as output.
 - Suppose, computer performs its task and comfortably passes the Turing test
 - Does the machine literally “understand” Chinese? – **Strong AI**
 - Is it merely simulating the ability to understand Chinese? – **Weak AI**

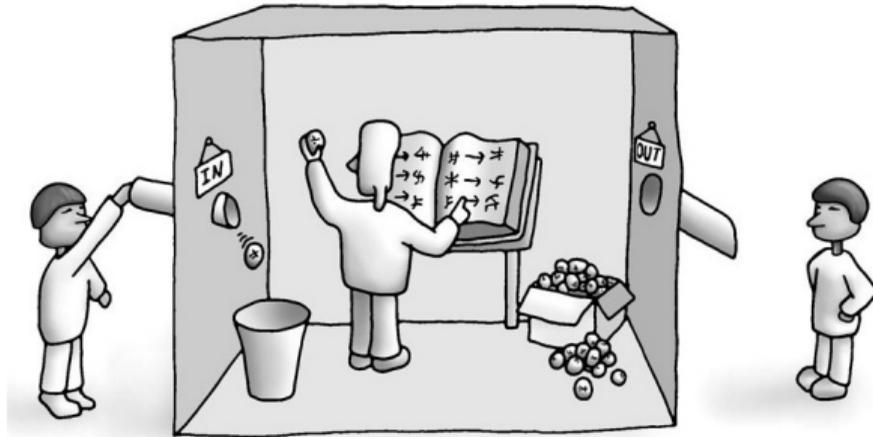


image source: Wikipedia

History of AI

- 1950: Alan Turing & The Turing Test - imitation Game
- 1956: The Dartmouth Conference - the birth of the field
- 1950s-70s: The era of great expectations
 - Logic Theorist (1955): The first AI program that proved mathematical theorems
 - ELIZA (1966): The first chatbot (simulate a psychotherapist)

ELIZA

Welcome to

```
EEEEEE LL      IIII  ZZZZZZ  AAAAA
EE      LL      II    ZZ     AA   AA
EEEEEE LL      II    ZZZ    AAAAAA
EE      LL      II    ZZ     AA   AA
EEEEEE LLLLLL  IIII  ZZZZZZ  AA   AA
```

Eliza is a mock Rogerian psychotherapist.

The original program was described by Joseph Weizenbaum in 1966.

This implementation by Norbert Landsteiner 2005.

ELIZA: Is something troubling you ?

YOU: Men are all alike.

ELIZA: What is the connection, do you suppose ?

YOU: They're always bugging us about something or other.

ELIZA: Can you think of a specific example ?

YOU: Well, my boyfriend made me come here.

ELIZA: Is it important to you that your boyfriend made you come here ?

YOU: He says I'm depressed much of the time.

ELIZA: I am sorry to hear that you are depressed.

YOU: It's true. I am unhappy.

ELIZA: Can you explain what made you unhappy ?

YOU:

Problem space

- Problems those can be solved easily
 - Finding maximum of n numbers
 - Problem can be described formally, computers can solve easily

Problem space

- Problems those can be solved easily
 - Finding maximum of n numbers
 - Problem can be described formally, computers can solve easily
- Problems that are intellectually challenging but *easy* for computers
 - Example - chess, traveling salesman problem, etc.
 - Problem can be described formally, computing requirements may be significant

Problem space

- Problems those can be solved easily
 - Finding maximum of n numbers
 - Problem can be described formally, computers can solve easily
- Problems that are intellectually challenging but *easy* for computers
 - Example - chess, traveling salesman problem, etc.
 - Problem can be described formally, computing requirements may be significant
- Problems that are easy for human being but *relatively tough* for computers
 - Example - object identification, speech-to-text conversion
 - Difficult to describe formally

Problem space

- Problems those can be solved easily
 - Finding maximum of n numbers
 - Problem can be described formally, computers can solve easily
- Problems that are intellectually challenging but *easy* for computers
 - Example - chess, traveling salesman problem, etc.
 - Problem can be described formally, computing requirements may be significant
- Problems that are easy for human being but *relatively tough* for computers
 - Example - object identification, speech-to-text conversion
 - Difficult to describe formally
- Problems that *cannot* be solved by computers
 - Example - detection of infinite loop in a program

Application domains

- Computer vision
- VLSI
- NLP
- Transportation
- Vehicle routing
- Airline/train scheduling
- Supply chain management
- Smart grid
- Operations research
- Healthcare
- Economics
- Games
- Robotics
- Weather forecast
- Agriculture
- Matching
- Linguistic
- Proving theorem
- etc.

Application domains

- Computer vision
- VLSI
- NLP
- Transportation
- Vehicle routing
- Airline/train scheduling
- Supply chain management
- Smart grid
- Operations research
- Healthcare
- Economics
- Games
- Robotics
- Weather forecast
- Agriculture
- Matching
- Linguistic
- Proving theorem
- etc.

Many problem domains but solved using a few set of techniques.

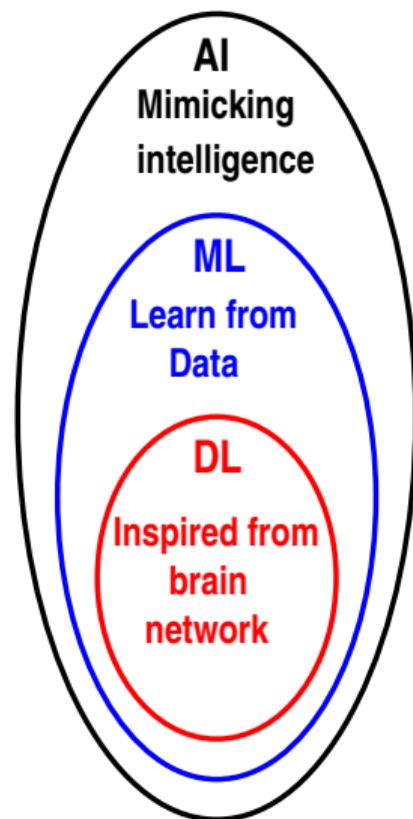
Broad techniques

- **Classical AI**
 - Search
 - Logic
 - Constraint satisfaction
 - Planning
 - Probabilistic reasoning

- **Statistical / Data driven AI**
 - Machine learning - symbolic / logical, statistical
 - Deep learning

Broad techniques

- **Classical AI**
 - Search
 - Logic
 - Constraint satisfaction
 - Planning
 - Probabilistic reasoning
- **Statistical / Data driven AI**
 - Machine learning - symbolic / logical, statistical
 - Deep learning



A search problem

- Let a factory produces N different products using M different components. Product n requires S_{nm} number of m th components. Each component has a cost C_m . There are limited number of components available (say L_m number of components for the m th item). Each product earns a revenue of C_n .
 - Given a set of demand for products, what is the best way to fulfill such demands?

A search problem

- Let a factory produces N different products using M different components. Product n requires S_{nm} number of m th components. Each component has a cost C_m . There are limited number of components available (say L_m number of components for the m th item). Each product earns a revenue of C_n .
 - Given a set of demand for products, what is the best way to fulfill such demands?
 - Product - white, milk, dark chocolate, Components - milk, butter, vanilla - Food industry

A search problem

- Let a factory produces N different products using M different components. Product n requires S_{nm} number of m th components. Each component has a cost C_m . There are limited number of components available (say L_m number of components for the m th item). Each product earns a revenue of C_n .
 - Given a set of demand for products, what is the best way to fulfill such demands?
 - Product - white, milk, dark chocolate, Components - milk, butter, vanilla - Food industry
 - Product - ring, pendant, earring, Components - gold, silver, diamond - Jewelry industry

A search problem

- Let a factory produces N different products using M different components. Product n requires S_{nm} number of m th components. Each component has a cost C_m . There are limited number of components available (say L_m number of components for the m th item). Each product earns a revenue of C_n .
 - Given a set of demand for products, what is the best way to fulfill such demands?
 - Product - white, milk, dark chocolate, Components - milk, butter, vanilla - Food industry
 - Product - ring, pendant, earring, Components - gold, silver, diamond - Jewelry industry
 - Product - sulfuric acid, aqua regia, Components - sulfur, chlorine, nitric acid - Chemical industry

Constraint satisfaction

- **Food nutrition** - Is it possible for a student to eat at several food joints at IITP and achieve the basic requirement of balanced diet? Each type of food has some basic nutrition value and maximum amount food intake by a student is limited.

Constraint satisfaction

- **Food nutrition** - Is it possible for a student to eat at several food joints at IITP and achieve the basic requirement of balanced diet? Each type of food has some basic nutrition value and maximum amount food intake by a student is limited.
- **Map coloring** - Given a world map and k number of colors, goal is to color the country such that no two countries that share the boundary have the same color.

Constraint satisfaction

- **Food nutrition** - Is it possible for a student to eat at several food joints at IITP and achieve the basic requirement of balanced diet? Each type of food has some basic nutrition value and maximum amount food intake by a student is limited.
- **Map coloring** - Given a world map and k number of colors, goal is to color the country such that no two countries that share the boundary have the same color.
- **Job-shop scheduling** - Consider car assembly consisting of the tasks: install axles (front and back), affix all four wheels (right and left, front and back), tighten nuts for each wheel, affix hubcaps, and inspect the final assembly. Each task requires some time, some tasks are dependent. Is it possible to finish the job within a given time T say?

Constraint satisfaction

- **Food nutrition** - Is it possible for a student to eat at several food joints at IITP and achieve the basic requirement of balanced diet? Each type of food has some basic nutrition value and maximum amount food intake by a student is limited.
- **Map coloring** - Given a world map and k number of colors, goal is to color the country such that no two countries that share the boundary have the same color.
- **Job-shop scheduling** - Consider car assembly consisting of the tasks: install axles (front and back), affix all four wheels (right and left, front and back), tighten nuts for each wheel, affix hubcaps, and inspect the final assembly. Each task requires some time, some tasks are dependent. Is it possible to finish the job within a given time T say?
- **Real life problems will have more number of constraints**

Logic

There are five houses.

The Englishman lives in the red house.

The Spaniard owns the dog.

Coffee is drunk in the green house.

The Ukrainian drinks tea.

The green house is immediately to the right of the ivory house.

The Old Gold smoker owns snails.

Kools are smoked in the yellow house.

Milk is drunk in the middle house.

The Norwegian lives in the first house.

The man who smokes Chesterfields lives in the house next to the man with the fox.

Kools are smoked in the house next to the house where the horse is kept.

The Lucky Strike smoker drinks orange juice.

The Japanese smokes Parliaments.

The Norwegian lives next to the blue house.

Logic

There are five houses.

The Englishman lives in the red house.

The Spaniard owns the dog.

Coffee is drunk in the green house.

The Ukrainian drinks tea.

The green house is immediately to the right of the ivory house.

The Old Gold smoker owns snails.

Kools are smoked in the yellow house.

Milk is drunk in the middle house.

The Norwegian lives in the first house.

The man who smokes Chesterfields lives in the house next to the man with the fox.

Kools are smoked in the house next to the house where the horse is kept.

The Lucky Strike smoker drinks orange juice.

The Japanese smokes Parliaments.

The Norwegian lives next to the blue house.

Who drinks water?
Who owns the zebra?

Logic

There are five houses.

The Englishman lives in the red house.

The Spaniard owns the dog.

Coffee is drunk in the green house.

The Ukrainian drinks tea.

The green house is immediately to the right of the ivory house.

The Old Gold smoker owns snails.

Kools are smoked in the yellow house.

Milk is drunk in the middle house.

The Norwegian lives in the first house.

The man who smokes Chesterfields lives in the house next to the man with the fox.

Kools are smoked in the house next to the house where the horse is kept.

The Lucky Strike smoker drinks orange juice.

The Japanese smokes Parliaments.

The Norwegian lives next to the blue house.

Who drinks water?
Who owns the zebra?

Zebra puzzle

Planning

- Three missionaries and three cannibals are on one side of a river that they wish to cross. A boat is available that can hold at most two people and at least one. You must never leave a group of missionaries outnumbered by cannibals on the same bank.

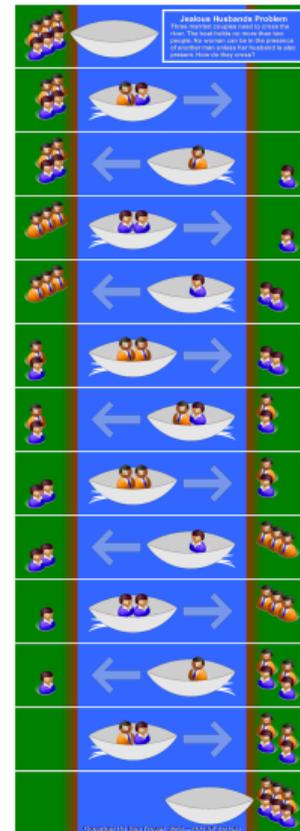


image source: Wikipedia

Planning

- Three missionaries and three cannibals are on one side of a river that they wish to cross. A boat is available that can hold at most two people and at least one. You must never leave a group of missionaries outnumbered by cannibals on the same bank.
- **Missionaries and cannibals problem, jealous husbands problem**

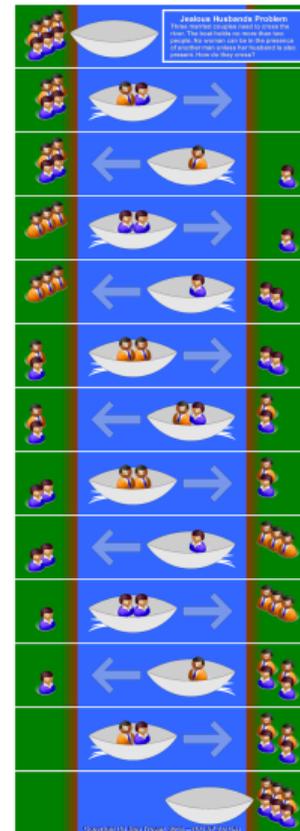


image source: Wikipedia

Planning

- Three missionaries and three cannibals are on one side of a river that they wish to cross. A boat is available that can hold at most two people and at least one. You must never leave a group of missionaries outnumbered by cannibals on the same bank.
- **Missionaries and cannibals problem, jealous husbands problem**
- **Applications - dialog management, risk management, robotics, etc.**

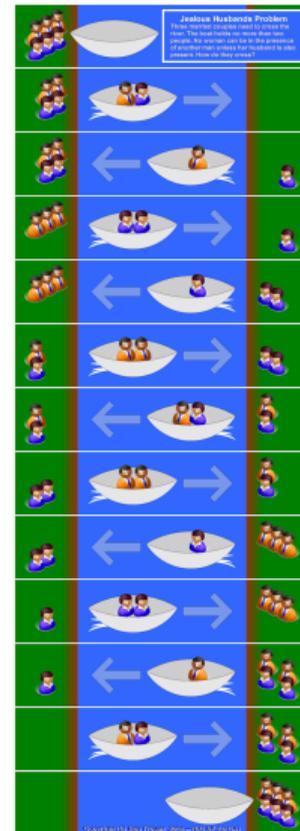
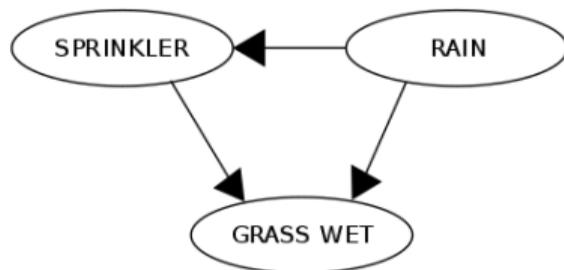


image source: Wikipedia

Probabilistic reasoning

- Two events can cause grass to be wet: an active sprinkler or rain. Rain has a direct effect on the use of the sprinkler
 - What is the probability that it would rain, given that we wet the grass?

RAIN	SPRINKLER	
	T	F
F	0.4	0.6
T	0.01	0.99



RAIN	T	F
	0.2	0.8

SPRINKLER	RAIN	GRASS WET	
		T	F
F	F	0.0	1.0
F	T	0.8	0.2
T	F	0.9	0.1
T	T	0.99	0.01

Real life applications of AI

- **AG** - Crop, soil monitoring, weather forecast

Real life applications of AI

- **AG** - Crop, soil monitoring, weather forecast
- **AR** - Smart city, urban planning

Real life applications of AI

- **AG** - Crop, soil monitoring, weather forecast
- **AR** - Smart city, urban planning
- **BT** - Drug discovery, diagnostics

Real life applications of AI

- **AG** - Crop, soil monitoring, weather forecast
- **AR** - Smart city, urban planning
- **BT** - Drug discovery, diagnostics
- **CH** - Organic chemical synthesis, structure elucidation

Real life applications of AI

- **AG** - Crop, soil monitoring, weather forecast
- **AR** - Smart city, urban planning
- **BT** - Drug discovery, diagnostics
- **CH** - Organic chemical synthesis, structure elucidation
- **EE** - Unit commitment, demand side management

Real life applications of AI

- **AG** - Crop, soil monitoring, weather forecast
- **AR** - Smart city, urban planning
- **BT** - Drug discovery, diagnostics
- **CH** - Organic chemical synthesis, structure elucidation
- **EE** - Unit commitment, demand side management
- **Finance** - Risk assessment/management, fraud detection

Real life applications of AI

- **AG** - Crop, soil monitoring, weather forecast
- **AR** - Smart city, urban planning
- **BT** - Drug discovery, diagnostics
- **CH** - Organic chemical synthesis, structure elucidation
- **EE** - Unit commitment, demand side management
- **Finance** - Risk assessment/management, fraud detection
- **IM** - Signal processing, smart sensors

Real life applications of AI

- **AG** - Crop, soil monitoring, weather forecast
- **AR** - Smart city, urban planning
- **BT** - Drug discovery, diagnostics
- **CH** - Organic chemical synthesis, structure elucidation
- **EE** - Unit commitment, demand side management
- **Finance** - Risk assessment/management, fraud detection
- **IM** - Signal processing, smart sensors
- **Law** - Logical analysis

Real life applications of AI

- **AG** - Crop, soil monitoring, weather forecast
- **AR** - Smart city, urban planning
- **BT** - Drug discovery, diagnostics
- **CH** - Organic chemical synthesis, structure elucidation
- **EE** - Unit commitment, demand side management
- **Finance** - Risk assessment/management, fraud detection
- **IM** - Signal processing, smart sensors
- **Law** - Logical analysis
- **ME** - Robotics

Real life applications of AI

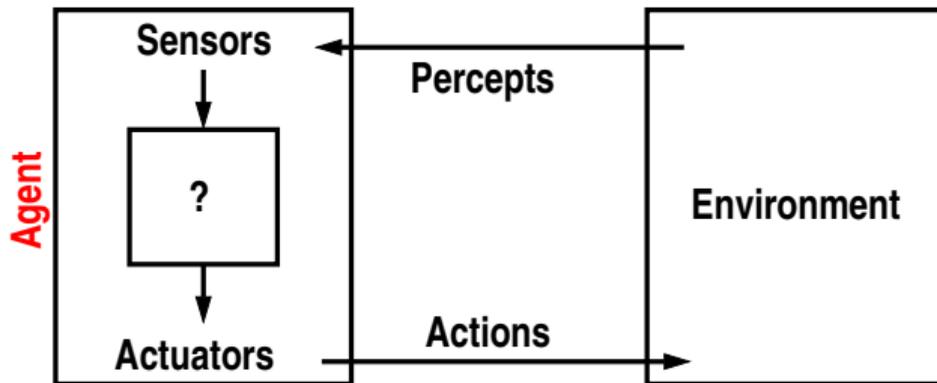
- **AG** - Crop, soil monitoring, weather forecast
- **AR** - Smart city, urban planning
- **BT** - Drug discovery, diagnostics
- **CH** - Organic chemical synthesis, structure elucidation
- **EE** - Unit commitment, demand side management
- **Finance** - Risk assessment/management, fraud detection
- **IM** - Signal processing, smart sensors
- **Law** - Logical analysis
- **ME** - Robotics
- **MA** - Proving theorem, formal verification

Real life applications of AI

- **AG** - Crop, soil monitoring, weather forecast
- **AR** - Smart city, urban planning
- **BT** - Drug discovery, diagnostics
- **CH** - Organic chemical synthesis, structure elucidation
- **EE** - Unit commitment, demand side management
- **Finance** - Risk assessment/management, fraud detection
- **IM** - Signal processing, smart sensors
- **Law** - Logical analysis
- **ME** - Robotics
- **MA** - Proving theorem, formal verification
- **OR** - Job shop scheduling, vehicle routing

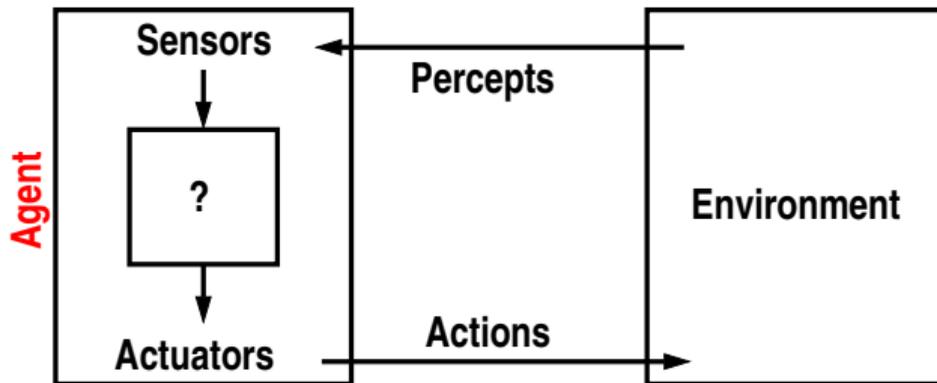
Agent and Environment

- An agent is anything that can be viewed as perceiving its environment through sensors and acting upon that environment through actuators
- Environment — Fully / partial observable, deterministic / stochastic, episodic / sequential, static / dynamic, discrete / continuous, known / unknown



Agent and Environment

- An agent is anything that can be viewed as perceiving its environment through sensors and acting upon that environment through actuators
- Environment — Fully / partial observable, deterministic / stochastic, episodic / sequential, static / dynamic, discrete / continuous, known / unknown
 - Known & partially observable - card games
 - Unknown & fully observable - video games



Thank you!