Introduction to Artificial Intelligence (G51IAI)

Dr Rong Qu

History of AI

Predictions in AI History

- AI
  - Originated in 1956, John McCarthy coined the term
  - very successful at early stage
- "Within 10 years a computer will be a chess champion"
  - Herbert Simon, 1957
  - IBM Deep Blue on 11 May 1997

Predictions in AI History

- Conversion from Russian to English
  - National Research Council, 1950s'
  - One example
    - "The spirit is willing but the flesh is weak" produced
    - "The vodka is good but the meat is rotten"
- Machine translations
  - Rendering the text from one language to another
  - Literal translation vs. free translation
  - Requires **knowledge** to establish content
  - Long way to go?

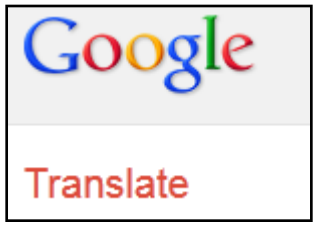

Why do we need AI anyway?

The Travelling Salesman Problem

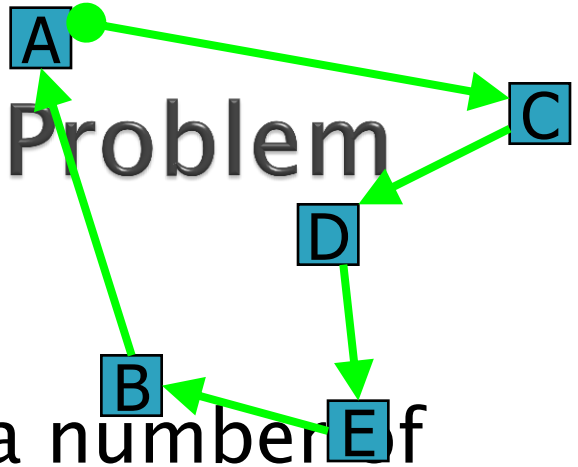

- A salesperson has to visit a number of cities
  - (S)He can start at any city and must finish at that same city
  - The salesperson must visit each city only once
- The number of possible routes is ??

The Travelling Salesman Problem

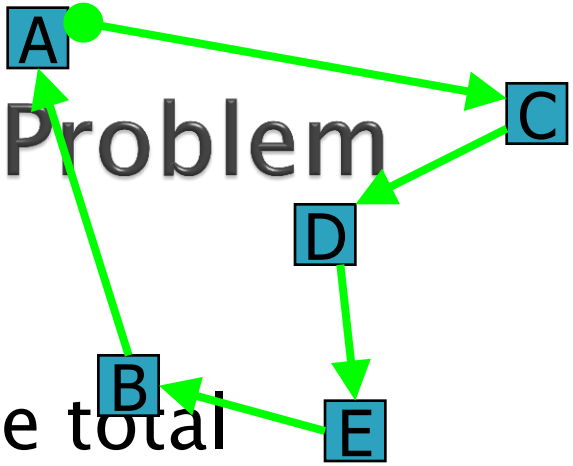

- The cost of a solution is the total distance traveled
- Solving the TSP means finding the minimum cost solution
  - Given a set of cities and distances between them
  - Find the optimal tour, i.e. the shortest possible such tour

Combinatorial Explosion

A 10 city TSP has 181,000 possible solutions

A 20 city TSP has 10,000,000,000,000,000 possible solutions

A 50 City TSP has 100,000,000,000,000,000,000,000,000,000,000,000,000,000,000,000,000,000,000,000,000,000 possible solutions

There are 1,000,000,000,000,000,000 litres of water on the planet*

*Mchalewicz, Z, Evolutionary Algorithms for Constrained Optimization Problems, CEC 2000 (Tutorial)

Combinatorial Explosion

A 50 City TSP has $1.52 * 10^{64}$ possible solutions

A 10GHz computer might do 10^9 tours per second

Running since start of universe, it would still only have done 10^{26} tours

Not even close to evaluating all tours!

One of the major unsolved theoretical problems in Computer Science

Combinatorial Explosion – Towers of Hanoi

- The original problem was stated that a group of Tibetan monks had to move 64 gold rings which were placed on diamond pegs

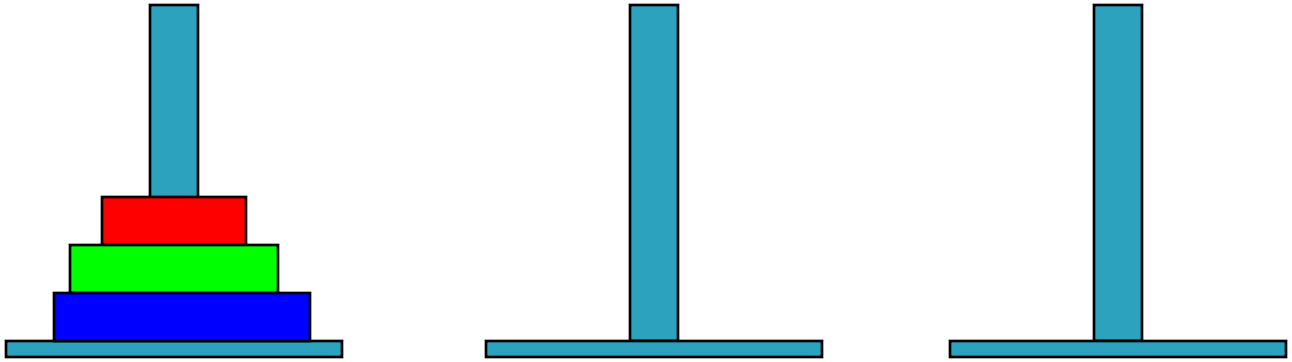

Combinatorial Explosion – Towers of Hanoi

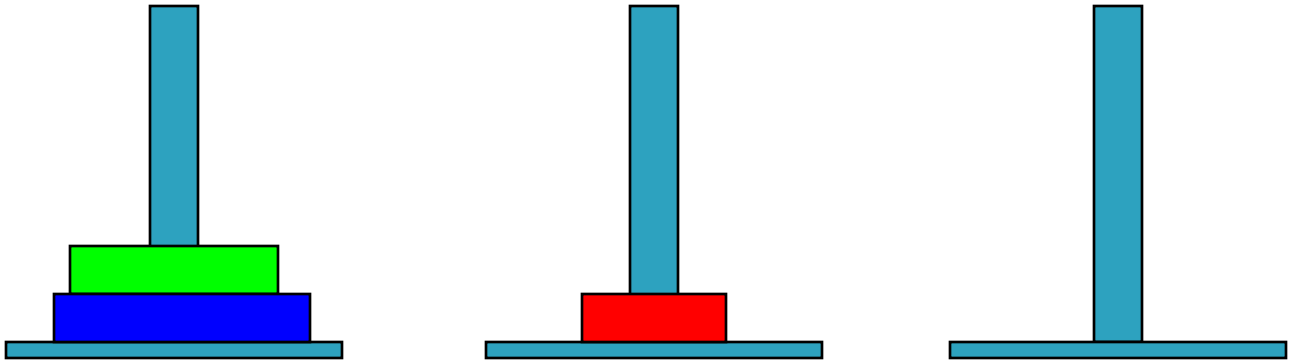

Combinatorial Explosion – Towers of Hanoi

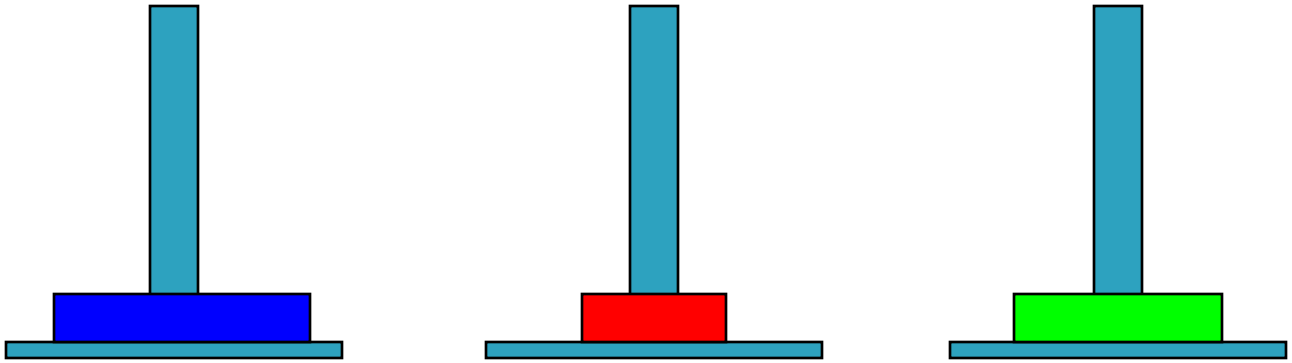

Combinatorial Explosion – Towers of Hanoi

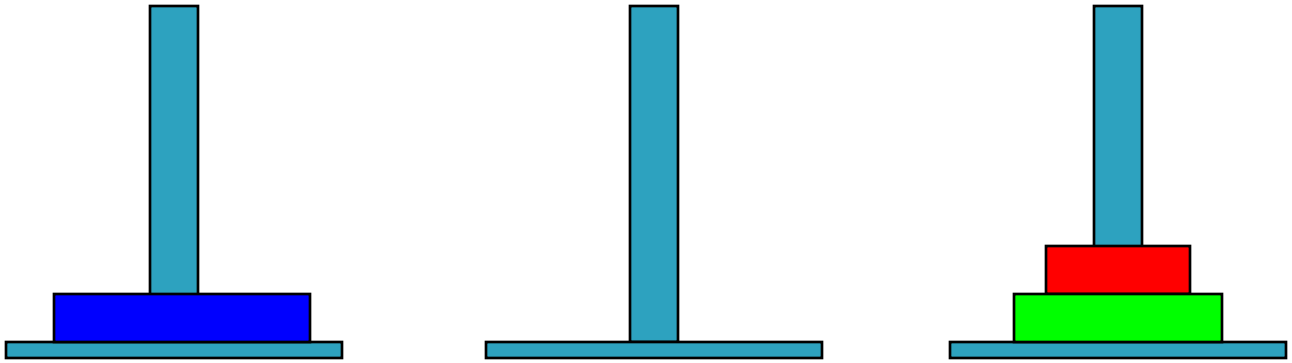

Combinatorial Explosion – Towers of Hanoi

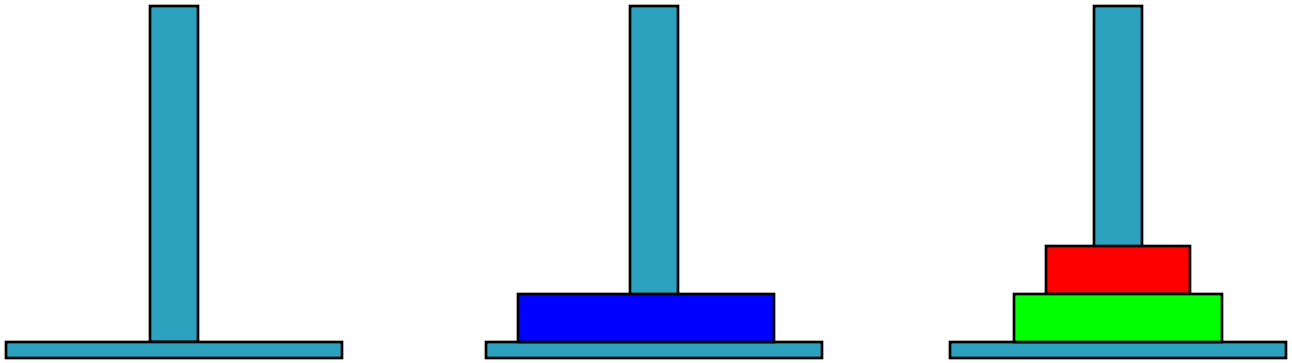

Combinatorial Explosion – Towers of Hanoi

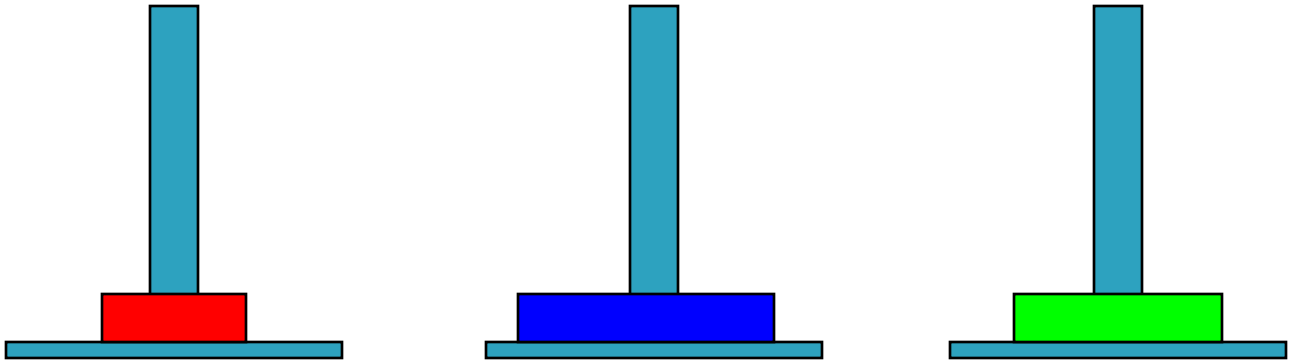

Combinatorial Explosion – Towers of Hanoi

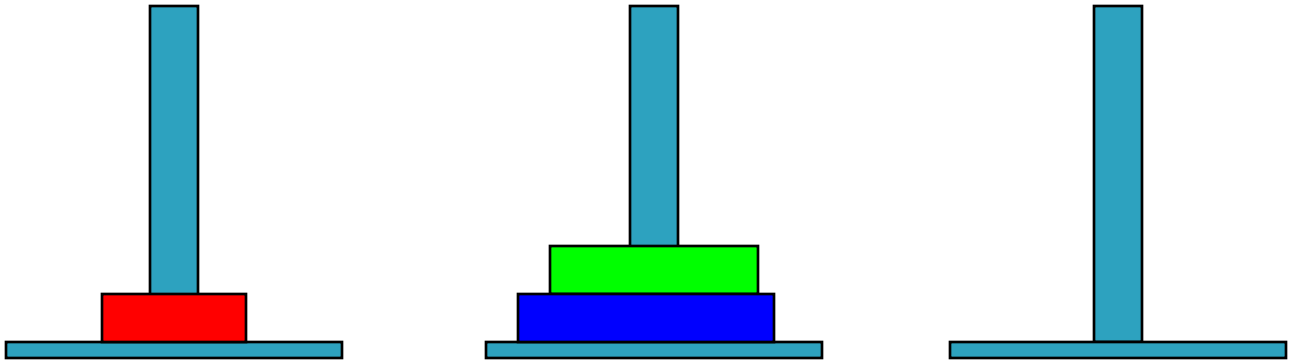

Combinatorial Explosion – Towers of Hanoi

- The original problem was stated that a group of Tibetan monks had to move 64 gold rings which were placed on diamond pegs

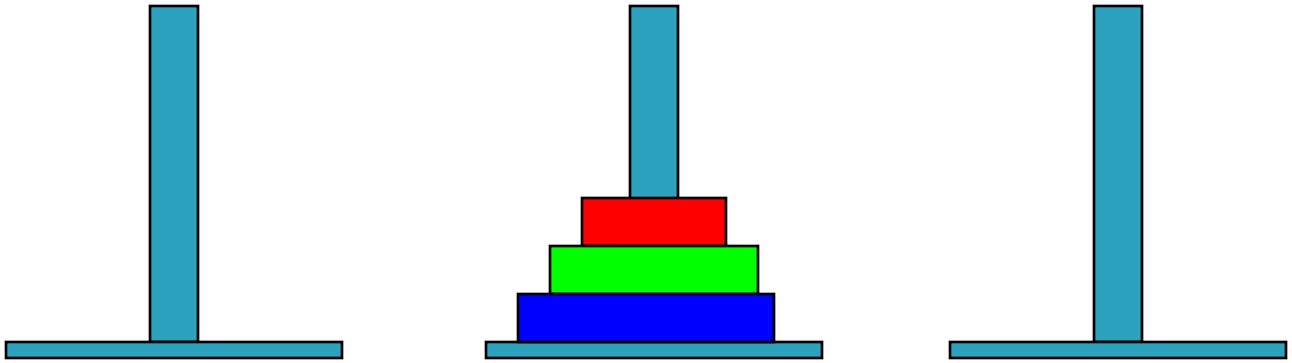

  - When they finished this task the world would end
  - Assume they could move one ring every second (or more realistically every five seconds)
  - How long till the end of the world?

Combinatorial Explosion – Towers of Hanoi

- It would require 3 trillion years!
- Using a computer we could do many more moves than one second, so go and try implementing the 64 rings towers of Hanoi problem
- If you are still alive at the end, try 1,000 rings!!!!

Combinatorial Explosion

- Optimize $f(x_1, x_2, \ldots, x_{100})$
  - where *f* is complex and x_i takes value of 0 or 1
  - The size of the **search space** is **?** $\cong 10^{30}$
- An exhaustive search is not an option
  - At 1,000 evaluations per second
  - Start the algorithm at the time the universe was created
  - As of now we would have considered just 1% of all possible solutions

Combinatorial Explosion

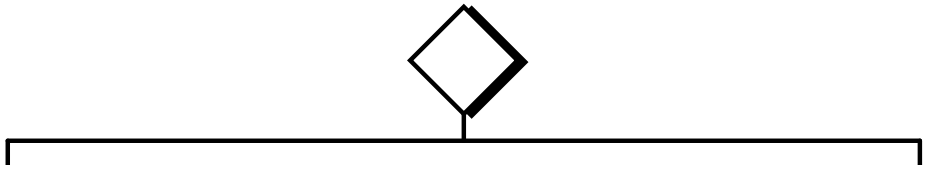

Combinatorial Explosion

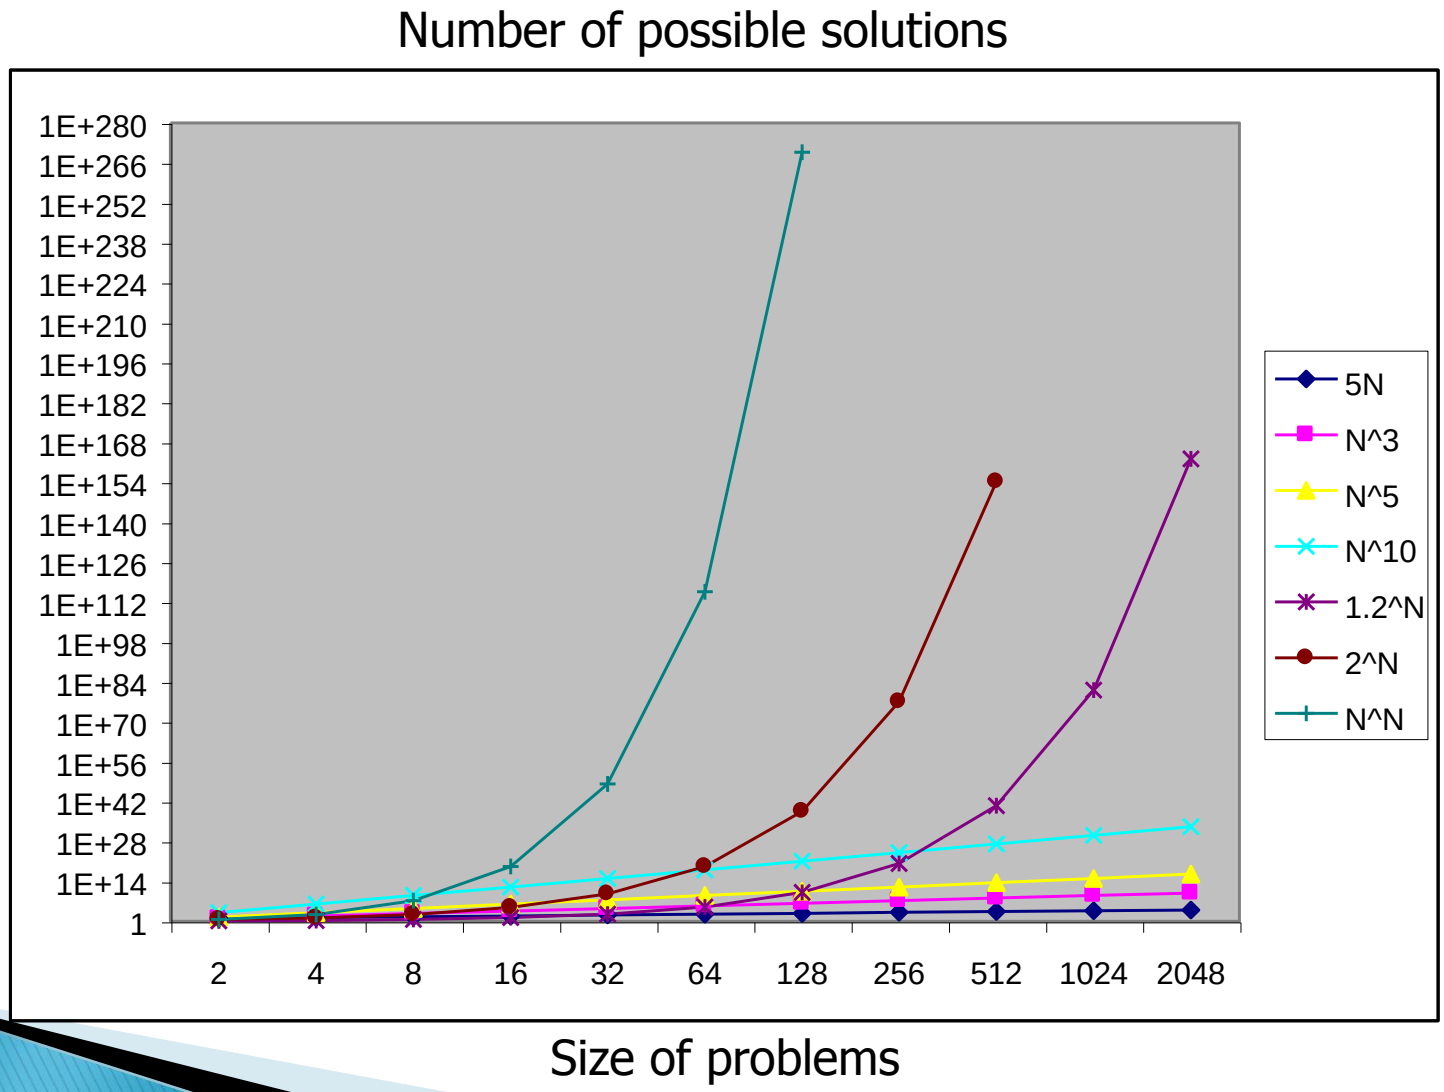

Combinatorial Explosion

the feature where the number of problem solutions grows **exponentially** with its size

Running on a computer capable of 1 million instructions/second

	10	20	50	100	200
N^2	1/10,000 second	1/2500 second	1/400 second	1/100 second	1/25 second
N^5	1/10 second	3.2 seconds	5.2 minutes	2.8 hours	3.7 days
2^N	1/1000 second	1 second	35.7 years	> 400 trillion centuries	45 digit no. of centuries
N^N	2.8 hours	3.3 trillion years	70 digit no. of centuries	185 digit no. of centuries	445 digit no. of centuries

Harel, D. 2000. Computer Ltd. : What they really can't do, Oxford University Press

Alan Turing

- Founder of computer science, mathematician, philosopher and code breaker
- Father of modern computer science

- **Turing test**

Alan Turing

- 1912 (23 June): Birth, Paddington, London
- 1931–34: Undergraduate at King's College, Cambridge University
- 1935: Elected fellow of King's College, Cambridge
- 1936: The Turing machine: *On Computable Numbers* Submitted
- 1936–38: At Princeton University. Ph.D. Papers in logic, algebra, number theory
- 1938–39: Return to Cambridge

Alan Turing

- 1939–40 Devises the Bombe, machine for Enigma decryption
- 1947–48: Papers on programming, neural nets, and prospects for artificial intelligence
- 1948: Manchester University
- 1950: Philosophical paper on machine intelligence: the Turing Test
- 1954 (7 June): Death by cyanide poisoning, Wilmslow, Cheshire

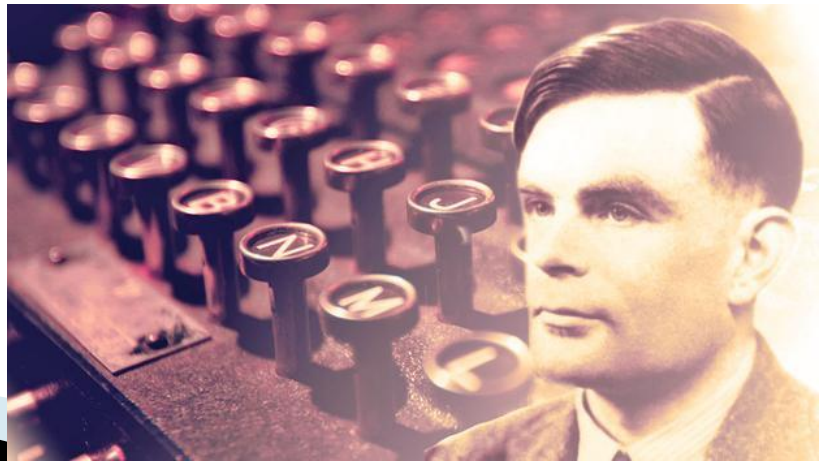

The Turing Test

- An interrogator is connected to a person and a machine via a terminal of some kind and cannot see either the person or machine.
- The interrogator's task is to find out which of the two candidates is the machine, and which is human, only by asking them questions
- If the machine can fool the interrogator 30% of the time, the machine is considered intelligent

The Turing Test

- Proposed by Alan Turing in 1950
  - Turing, A.M. 1950. "Computing Machinery and Intelligence." Mind LIX, no. 2236, 1950 : 433-460
  - http://www.loebner.net/Prizef/TuringArticle.html
- If the Turing Test was passed Turing would conclude that the machine was intelligent
  - The ELIZA program and Internet chatbot MGONZ have fooled humans
  - The A.L.I.C.E. program fooled one judge in the 2001 Loebner Prize Competition
- Suggested as a way of saying when we could consider machines to be **intelligent**

The Turing Test

- Question : “What is 35,076 divided by 4,567?”

  Answer : ????

  Answer : 7.6803153

http://cogsci.ucsd.edu/~asaygin/tt/ttest.html

The Turing Test

- You're on the internet chatting to two others "A" and "B"
  - one is a person
  - one is a machine trying to imitate a person (e.g. capable of discussing the X-factor?)
- If you can't tell the difference
  - then the machine must be intelligent
  - Or at least **act** intelligent?

The Turing Test

- The computer needs (at least) the following capabilities:
  - Knowledge representation
  - Automated reasoning
  - Machine learning
  - Computer vision
  - Robotics

The Turing Test

- Does this test show intelligence?
  - How about the person doesn't know x-factor
  - If the person doesn't speak English?
- Is this test about
  - Behaviour

  or
  - Intelligence
- Provides a satisfactory **operational** definition of intelligence

The Chinese Room Experiment

- In 1980 John Searle devised a thought experiment which he called the Chinese Room
  - *Searle, J.R. 1980. Minds, brains and programs. Behavioural and Brain Sciences, 3: 417-457, 1980
  - **Behaving** intelligently was not enough to prove a computer was intelligent

*http://members.aol.com/NeoNoetics/MindsBrainsPrograms.html

The Chinese Room Experiment

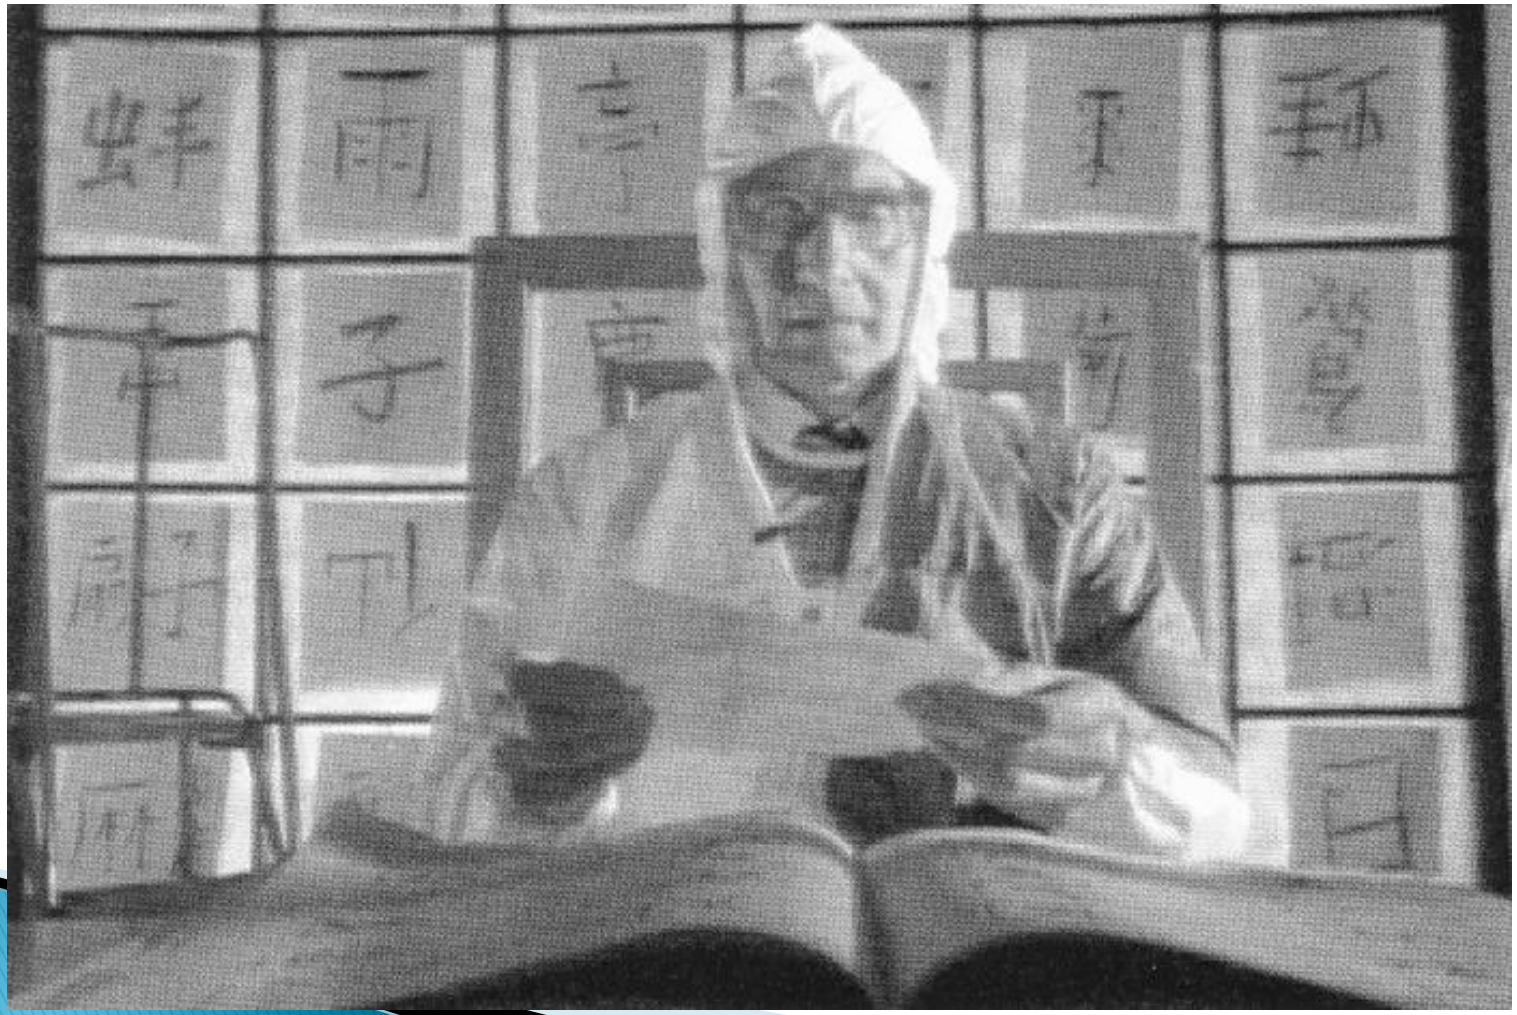

The Chinese Room Experiment

- Simple Rule processing system
  - in which the "rule processor" happens to be intelligent
  - but has no **understanding** of the rules
- Does the system understand Chinese?
  - Just comprises a rule book and papers
- But the system as a whole does **act** as it understands Chinese!
  - Regarded as invalid by many scientists
- Does Google Translate understand Chinese?

The Turing Test vs. Chinese Room

- You need to be able to
  - tell the differences of the objectives of these two tests
  - have an opinion about The **Turing Test** and **Chinese Room**

Summary

- Understand what is meant by **combinatorial explosion** (esp. wrt TSP)
- The **Turing Test** and **Chinese Room**
  - Be able to recognize the relationship between The Turing Test and The Chinese Room
- **Definitions of AI, strong vs. weak AI**

Self Study

- Read the following AIMA book chapters and understand
  - **AI terminologies**
    - Weak AI: Machine can possibly *act* intelligently
    - Strong AI: Machines can actually *think* intelligently
    - 4 categories of definitions from different AI books
  - **Turing test** (AIMA section 26.1)
    - A satisfactory operational definition of intelligence
  - The **Chinese Room** experiment (AIMA section 26.2)

Self Study

- Systems that

Think like *humans*

Act like *humans*

Think *rationally*

Act *rationally*

Lab Session

- You only need to attend one of these sessions
  - Friday 20th March, 11am-1pm
  - Friday 20th March, 3-5pm
  - Tuesday 24th March, 12-2pm
  - Thursday 26th March, 2-4pm

  - Monday 23rd March, 3-5pm (backup)
  - Friday 27th March, 3-5pm (backup)

  - Build ANNs in Matlab
    - Optional, but useful