BOSTON UNIVERSITY

CS640: Artificial Intelligence

Margrit Betke

1st lecture

BOSTON UNIVERSITY

Welcome to AI !

Prof. Margrit Betke

Beethoven

e e

Instructors: Yiwen Gu and Mahir Patel

TF: Stan (Sha) Lai

What is AI?

BOSTON UNIVERSITY

What is AI?

AI studies

- how to build "intelligent computers"
- how to make machines that exhibit characteristics associated with intelligence in humans

BOSTON UNIVERSITY

What is AI?

AI studies

- how to build "intelligent computers"
- how to make machines that exhibit characteristics associated with intelligence in humans

Machine that
do things that
would require
intelligence if done
by humans

BOSTON UNIVERSITY

What is AI?

AI studies

- how to build "intelligent computers"
- how to make machines that exhibit characteristics associated with intelligence in humans

- think, reason
- solve problems
- learn
- understand language

Machine that
do things that
would require
intelligence if done
by humans

BOSTON UNIVERSITY

"Modern" View of AI:

AI studies computations for

- perception
- reasoning ← "classic AI"
- action

AI is

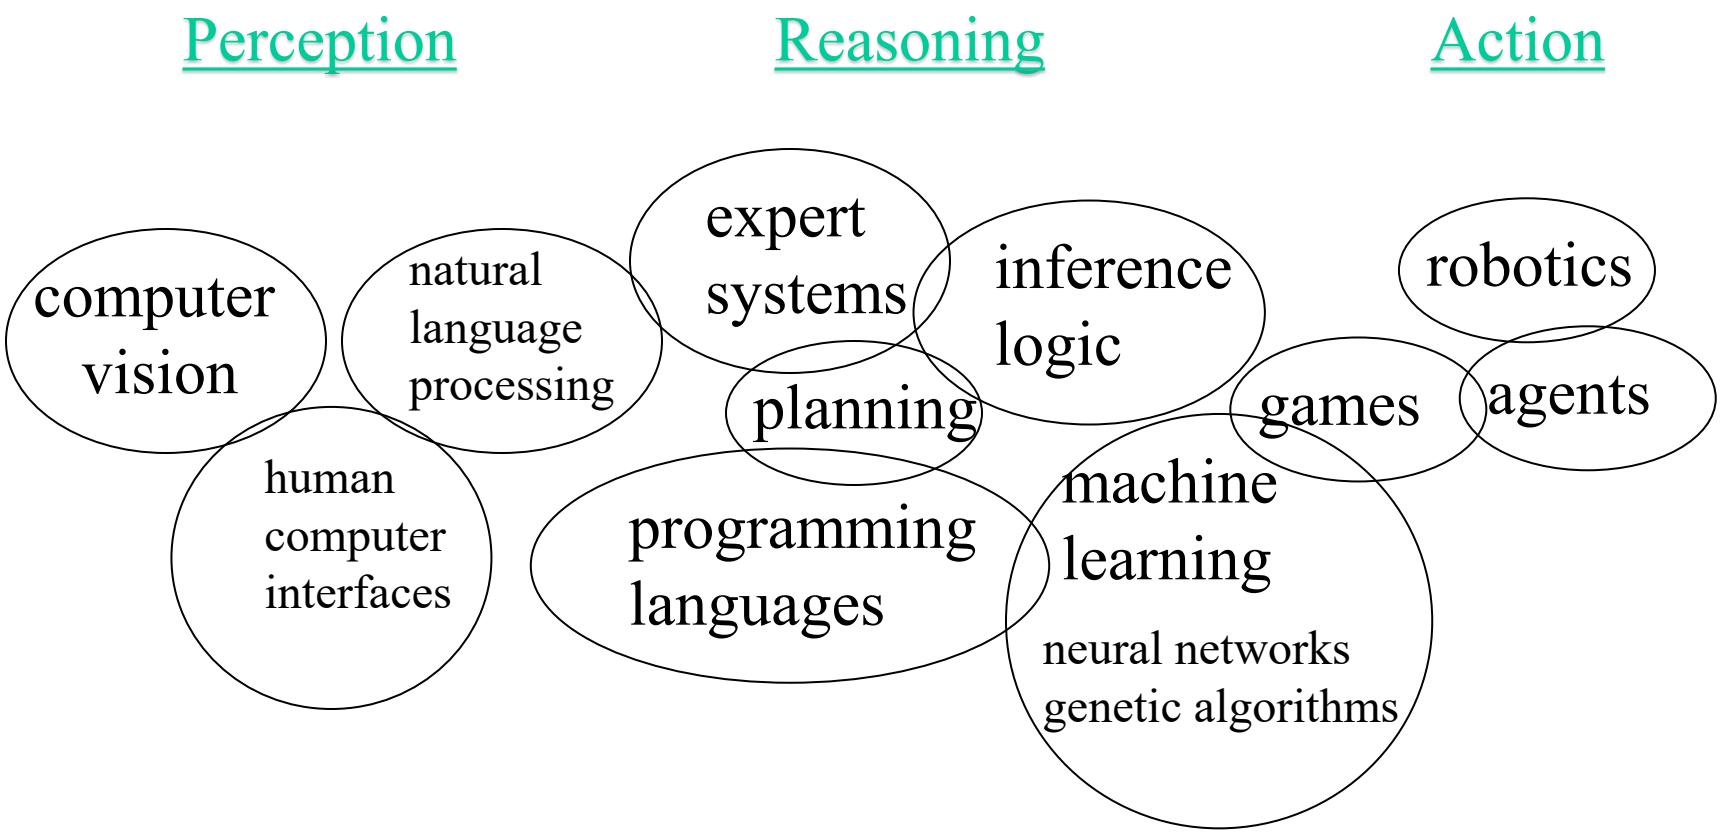
Perception
Reasoning
Action
computer vision
natural language processing
human computer interfaces
expert systems
inference logic
planning
programming languages
machine learning
neural networks
genetic algorithms
games
robotics
agents

BOSTON UNIVERSITY

Agent-oriented View of AI:

An Agent

- is (semi-) autonomous
- does independent piece of problem solving
- is "situated," i.e., sensitive to its own environment
- belongs to society of agents and interacts with other agents

Intelligence emerges from society of agents

BOSTON UNIVERSITY

Alan Turing's Test

"Can machines pass a behavior test for intelligence?"

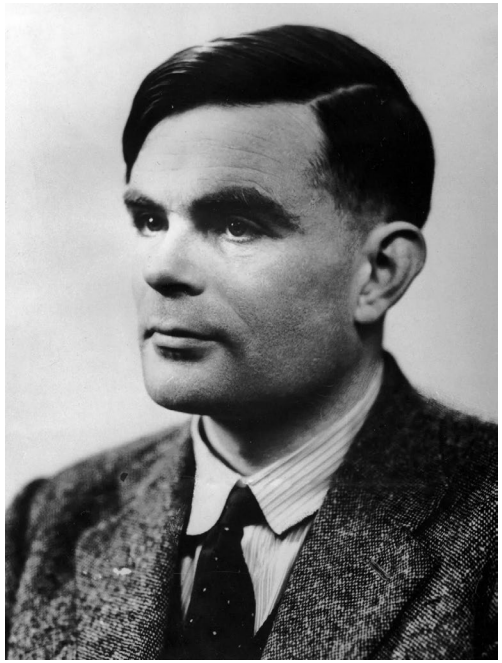

BOSTON UNIVERSITY

Alan Turing's Test = Imitation Game

"Can machines pass a behavior test for intelligence?"

Alan Turing's Test = Imitation Game

"Can machines pass a behavior test for intelligence?"

Person: "Are you the woman?"

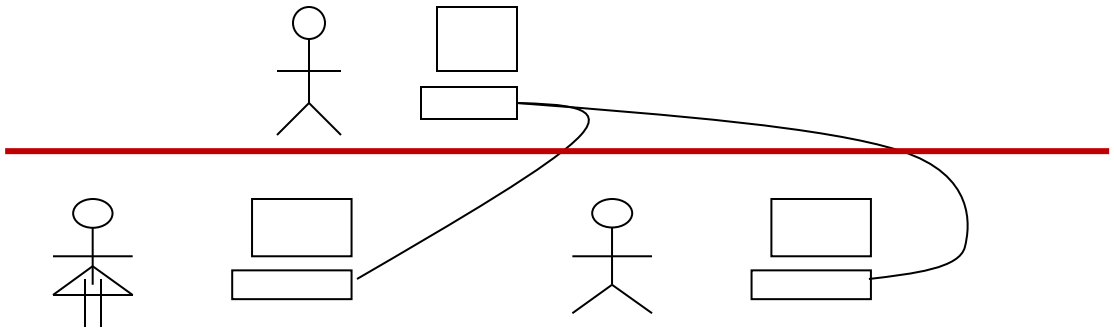

Can person tell the difference?

She: "I'm the woman." He: "I'm the woman."

Alan Turing's Test = Imitation Game

"Can machines pass a behavior test for intelligence?"

Person: "Are you the woman?"

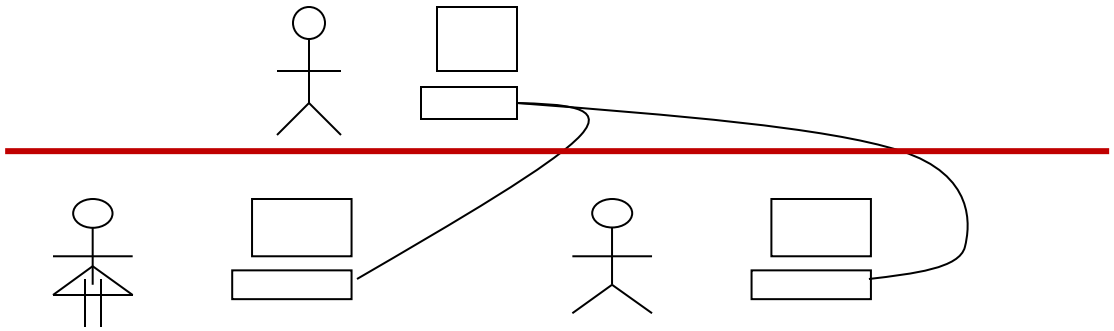

Can person tell the difference?

She: "I'm the woman."

Phase 1: He: "I'm the woman."

Phase 2: Computer : "I'm the woman."

BOSTON UNIVERSITY

Turing's Prediction (1950):

In 2000, a computer will have a X% chance of deceiving a human interrogator that it was human in a Y min conversation.

What do you think is X? What Y?

BOSTON UNIVERSITY

Turing's Prediction (1950):

In 2000, a computer will have a 30% chance of deceiving a human interrogator that it was human in a 5 min conversation.

BOSTON UNIVERSITY

Turing's Prediction (1950):

In 2000, a computer will have a 30% chance of deceiving a human interrogator that it was human in a 5 min conversation.

How about now? 2024

Large Language Models like ChatGPT?

BOSTON UNIVERSITY

Turing's Prediction (1950):

In 2000, a computer will have a 30% chance of deceiving a human interrogator that it was human in a 5 min conversation.

How about now? 2024

Large Language Models like ChatGPT?

Not good at understanding logical inference, moving from premises to logical consequences

BOSTON UNIVERSITY

LLMs are not good at understanding logical inference, moving from premises to logical consequences

Prof. Leonid Levin's story about how his PhD advisor Andrey Kolmogorov saved Soviet mathematicians.

A: The village priest asks God for rain.

B: It rains the next day.

Does A entail B?

A: The village priest asks God for rain.

B: It rains the next day.

Does A entail B? No. These are independent statements.

B': The priest prays to God.

Does A entail B'? Yes.

A: The village priest asks God for rain.

B: It rains the next day.

Does A entail B? No. These are independent statements.

B': The priest prays to God.

Does A entail B'? Yes.

B'': The village priest is unconscious.

Does A entail B''?

A: The village priest asks God for rain.

B: It rains the next day.

Does A entail B? No. These are independent (neutral) statements.

B': The priest prays to God.

Does A entail B'? Yes.

B'': The village priest is unconscious.

Does A entail B''? No. This is a contradiction.

BOSTON UNIVERSITY

Joseph Weizenbaum's Eliza (1966):

Interactive program that mimics a psychologist.

Goal: De-mystify computers

Results: lots of misunderstandings

concern for "social implications of computers"

Emacs version of Eliza in action:

M-x doctor

Web version:

https://www.cyberpsych.org/eliza

BOSTON UNIVERSITY

Learning Objectives

Be able to:

- Define AI in various ways
- Explain the Turing Test
- Define the task of inference in NLP and provide examples
- Know about the first “therapist” chatbot Eliza

BOSTON UNIVERSITY

Next: Going through course syllabus

http://www.cs.bu.edu/faculty/betke/cs640

Then:

http://www.cs.bu.edu/faculty/betke/cs640/open/AI-successes-lecture.html